

Docket: 43876-155

PATENT

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

| | | |
|--|---|---------------------------|
| In re Application of | : | Customer Number: 20277 |
| | : | |
| HANSEN et al. | : | Confirmation Number: 4560 |
| | : | |
| Application No.: 10/757,516 | : | Group Art Unit: 2183 |
| | : | |
| Filed: January 15, 2004 | : | Examiner: Eric Coleman |
| | : | |
| For: PROGRAMMABLE PROCESSOR AND SYSTEM FOR STORE MULTIPLEX OPERATION | | |

SUPPLEMENTAL DECLARATION OF KORBIN VAN DYKE

I, Korbin Van Dyke, state that:

Summary of My Opinions

1. I previously submitted a declaration in connection with this proceeding (see Van Dyke Declaration dated August 15, 2007), referenced hereinafter as the "initial Van Dyke declaration". For brevity, I will not repeat information set forth in the initial Van Dyke declaration in this declaration.

2. In preparation of this declaration I have reviewed U.S. Patent Application Serial No. 10/757,516. I have also reviewed U.S. Patent Nos. 5,742,840 and 6,295,599 (respectively the '840 and '599 patents) that the 10/757,516 patent application indirectly claims priority to, as well as appendices to the '840 and '599 patents (the Terpsichore and Zeus System Architecture manuals, respectively, and hereinafter referred to respectively as the Terpsichore and the Zeus manuals). I have reviewed the Office Action for the 10/757,516 patent application mailed on October 30, 2007, including the paragraph on page 11 that discusses the Response to Arguments and particularly the Examiner's conclusion that the priority for the claimed invention does not extend to the '840 or the '599 patents, since limitations of claims 9, 18, 24, 25, 31, and 32 are not supported by the '840 or the '599 patents. My understanding is that the features of the claimed invention are taught and supported by complying with the written description requirement and the enablement requirement. My understanding of the written description requirement is that a patent disclosure must describe the claimed invention in sufficient detail that one of ordinary

Supplemental Declaration of Korbin S Van Dyke

skill in the art can reasonably conclude that the inventor had possession of the claimed invention at the time of filing the patent disclosure. My understanding of the enablement requirement is that the patent disclosure must contain sufficient information regarding the subject matter of the claims to enable one of ordinary skill in the pertinent art to make and use the claimed invention. I further understand that whether the enablement requirement is met depends on whether undue experimentation is necessary for one of skill in the art to practice the invention in light of the patent disclosure.

3. Based on my review of the materials identified in paragraph 2 of this declaration, it is my opinion that with respect to the following limitations relating to claims 9, 18, 24, 25, 31, and 32 (as amended), the disclosures of the '840 patent and the '599 patent each indicate that the inventors were in possession of the claimed invention of the 10/757,516 patent application as of the August 16, 1995 filing date of the '840 patent and further as of the August 24, 1999 filing date of the '599 patent; and further the disclosures of the '840 patent and the '599 patent each would have enabled a person of ordinary skill in the art to make and use, without undue experimentation, the claimed invention of the 10/757,516 patent application as of the August 16, 1995 filing date of the '840 patent, and further as of the August 24, 1999 filing date of the '599 patent. The limitations referred to are:

{ claim 9 } “wherein the execution unit is further operable to, in response to decoding a second single instruction specifying a register containing a first plurality of floating-point operands and another register containing a second plurality of floating-point operands, multiply the first plurality of floating-point operands by the second plurality of floating-point operands to produce a plurality of products and provide the plurality of products to partitioned fields of a result register as a catenated result”

{ claim 18 } “wherein the execution unit is further operable to, in response to decoding a second single instruction specifying a register containing a first plurality of floating-point operands and another register containing a second plurality of floating-point operands, multiply the first plurality of floating-point operands by the second plurality of floating-point operands to produce a plurality of products and provide the plurality of products to partitioned fields of a result register as a catenated result”

{ claim 24 } “wherein each of the first and second operands has a width of 64 bits”

{ claim 25 } “wherein the execution unit is further capable of executing a plurality of different group floating-point arithmetic operations that arithmetically operate on multiple floating-point operands stored in partitioned fields of registers in the register file to produce a catenated result that is returned to a register in the register file, wherein the catenated result comprises a plurality of individual floating-point results”

{ claim 31 } “wherein each of the first and second operands has a width of 64 bits”

{claim 32} “wherein the execution unit is further capable of executing a plurality of different group floating-point arithmetic operations that arithmetically operate on multiple floating-point operands stored in partitioned fields of registers in the register file to produce a catenated result that is returned to a register in the register file, wherein the catenated result comprises a plurality of individual floating-point results”

Summary of ‘840 Analysis:

4. The disclosure of the ‘840 patent provides detailed information and description that I believe indicates that the inventors were in possession of the aforementioned limitations of claim 9 (as amended) of the 10/757,516 patent application, and that I further believe would have enabled a person of ordinary skill in the art to make and use the claimed invention without undue experimentation. For example, on at least pages 19-21 (describing floating-point data formats), 24-25 (describing general registers), 29 and 47-48 (describing floating-point arithmetic hardware), and 129-131 of the Terpsichore manual (describing details of Group Floating-point instructions such as various forms of Group Floating-point Multiply instructions) there are detailed descriptions of the aforementioned claim elements.

5. The aforementioned limitations of claim 18 (a data processing system claim) are substantially similar to the aforementioned limitations of claim 9 (a programmable processor claim). Further, parent claim context providing antecedent basis for claim 18 (specifically “the execution unit”) is substantially similar to corresponding parent claim context of claim 9. Thus, for at least the reasons described in paragraph 4 of this declaration, the disclosure of the ‘840 patent provides detailed information and description that I believe indicates that the inventors were in possession of the aforementioned limitations of claim 18 (as amended) of the 10/757,516 patent application, and that I further believe would have enabled a person of ordinary skill in the art to make and use the claimed invention without undue experimentation.

6. The disclosure of the ‘840 patent provides detailed information and description that I believe indicates that the inventors were in possession of the aforementioned limitations of claim 24 (as amended) of the 10/757,516 patent application, and that I further believe would have enabled a person of ordinary skill in the art to make and use the claimed invention without undue experimentation. For example, on at least pages 24-25 (describing general registers), 26 (generally describing store instructions), and 150-157 of the Terpsichore manual (describing details of Store and Store Immediate instructions such as various forms of Store Immediate and Store Multiplex Immediate instructions) there are detailed descriptions of the aforementioned claim elements.

7. The disclosure of the ‘840 patent provides detailed information and description that I believe indicates that the inventors were in possession of the aforementioned limitations of claim 25 (as amended) of the 10/757,516 patent application, and that I further believe would have enabled a person of ordinary skill in the art to make and use the claimed invention without undue experimentation. For example, on at least pages 19-21 (describing floating-point data formats), 24-25 (describing general registers), 29 and 47-48 (describing floating-point arithmetic hardware), and 129-131 of the Terpsichore manual (describing details of Group Floating-point

Supplemental Declaration of Korbin S Van Dyke

instructions such as various Group Floating-point Add, Divide, and Multiply forms) there are detailed descriptions of the aforementioned claim elements.

8. The aforementioned limitations of claim 31 (a device claim) are substantially similar to the aforementioned limitations of claim 24 (a programmable processor claim). Further, parent claim context providing antecedent basis for claim 31 (specifically “the first and second operands”) is substantially similar to corresponding parent claim context of claim 24. Thus, for at least the reasons described in paragraph 6 of this declaration, the disclosure of the ‘840 patent provides detailed information and description that I believe indicates that the inventors were in possession of the aforementioned limitations of claim 31 (as amended) of the 10/757,516 patent application, and that I further believe would have enabled a person of ordinary skill in the art to make and use the claimed invention without undue experimentation.

9. The aforementioned limitations of claim 32 (a device claim) are substantially similar to the aforementioned limitations of claim 25 (a programmable processor claim). Further, parent claim context providing antecedent basis for claim 32 (specifically “the execution unit”) is substantially similar to corresponding parent claim context of claim 25. Thus, for at least the reasons described in paragraph 7 of this declaration, the disclosure of the ‘840 patent provides detailed information and description that I believe indicates that the inventors were in possession of the aforementioned limitations of claim 32 (as amended) of the 10/757,516 patent application, and that I further believe would have enabled a person of ordinary skill in the art to make and use the claimed invention without undue experimentation.

Summary of ‘599 Analysis:

10. The disclosure of the ‘599 patent provides detailed information and description that I believe indicates that the inventors were in possession of the aforementioned limitations of claim 9 (as amended) of the 10/757,516 patent application, and that I further believe would have enabled a person of ordinary skill in the art to make and use the claimed invention without undue experimentation. For example, on at least pages 14-16 (describing floating-point data formats), 19-20 (describing general registers), 23-24 and 55 (describing floating-point arithmetic hardware), and 258-260 of the Zeus manual (describing details of Ensemble Floating-point instructions such as various forms of Ensemble Multiply Floating-point instructions) there are detailed descriptions of the aforementioned claim elements. Note that in the Zeus manual, group floating-point instructions are termed “ensemble” floating-point instructions.

11. The aforementioned limitations of claim 18 are substantially similar to the aforementioned limitations of claim 9. Thus, for at least the reasons described in paragraph 10 of this declaration, the disclosure of the ‘599 patent provides detailed information and description that I believe indicates that the inventors were in possession of the aforementioned limitations of claim 18 (as amended) of the 10/757,516 patent application, and that I further believe would have enabled a person of ordinary skill in the art to make and use the claimed invention without undue experimentation.

12. The disclosure of the ‘599 patent provides detailed information and description that I believe indicates that the inventors were in possession of the aforementioned limitations of claim 24 (as amended) of the 10/757,516 patent application, and that I further believe would

Supplemental Declaration of Korbin S Van Dyke

have enabled a person of ordinary skill in the art to make and use the claimed invention without undue experimentation. For example, on at least pages 19-20 (describing general registers), 21 (generally describing store instructions), 123-125, and 128-130 of the Zeus manual (describing details of Store and Store Immediate instructions, including Store Multiplex and Store Multiplex Immediate forms) there are detailed descriptions of the aforementioned claim elements.

13. The disclosure of the '599 patent provides detailed information and description that I believe indicates that the inventors were in possession of the aforementioned limitations of claim 25 (as amended) of the 10/757,516 patent application, and that I further believe would have enabled a person of ordinary skill in the art to make and use the claimed invention without undue experimentation. For example, on at least pages 14-16 (describing floating-point data formats), 19-20 (describing general registers), 23-24 and 55 (describing floating-point arithmetic hardware), and 258-260 of the Zeus manual (describing details of Ensemble Floating-point instructions, such as various Ensemble Multiply, Add, and Divide forms) there are detailed descriptions of the aforementioned claim elements.

14. The aforementioned limitations of claim 31 are substantially similar to the aforementioned limitations of claim 24. Thus, for at least the reasons described in paragraph 12 of this declaration, the disclosure of the '599 patent provides detailed information and description that I believe indicates that the inventors were in possession of the aforementioned limitations of claim 31 (as amended) of the 10/757,516 patent application, and that I further believe would have enabled a person of ordinary skill in the art to make and use the claimed invention without undue experimentation.

15. The aforementioned limitations of claim 32 are substantially similar to the aforementioned limitations of claim 25. Thus, for at least the reasons described in paragraph 13 of this declaration, the disclosure of the '599 patent provides detailed information and description that I believe indicates that the inventors were in possession of the aforementioned limitations of claim 32 (as amended) of the 10/757,516 patent application, and that I further believe would have enabled a person of ordinary skill in the art to make and use the claimed invention without undue experimentation.

16. A detailed explanation of the basis for my opinions is set forth in the remainder of this declaration.

Detailed Basis for My Opinions

Analysis of the disclosures of the '840 and the '599 patents:

17. For brevity, the following analysis focuses on and provides details relating to the '840 patent, while reciting summary information pointing out where similar descriptive information is provided in the '599 patent.

18. As discussed in paragraph 20 of the initial Van Dyke declaration, the '840 patent describes structure of a general purpose, programmable media processor (including, for example, a register file and an execution unit), and the '840 patent recites that an instruction set for the

Supplemental Declaration of Korbin S Van Dyke

general purpose media processor is described by the Microfiche Appendix (referred to herein as the Terpsichore manual). In addition to elements discussed in paragraph 20 of the initial Van Dyke declaration, the '840 patent also describes that a unified stream of media data is processed by storage into the register file 110, and multi-precision arithmetic operations are performed on the media data. The operations include Boolean, integer, and floating-point mathematical operations (see '840, column 5, lines 47-53). Floating-point addition, subtraction, multiplication, division, and square root are supported in hardware ('840, column 15, lines 57-59). Similarly, the '599 patent describes structure of a general purpose, programmable processor for broadband applications, and the '599 patent includes and refers to a Microfiche Appendix (referred to herein as the Zeus manual) that describes, for example, an instruction set for the general purpose processor.

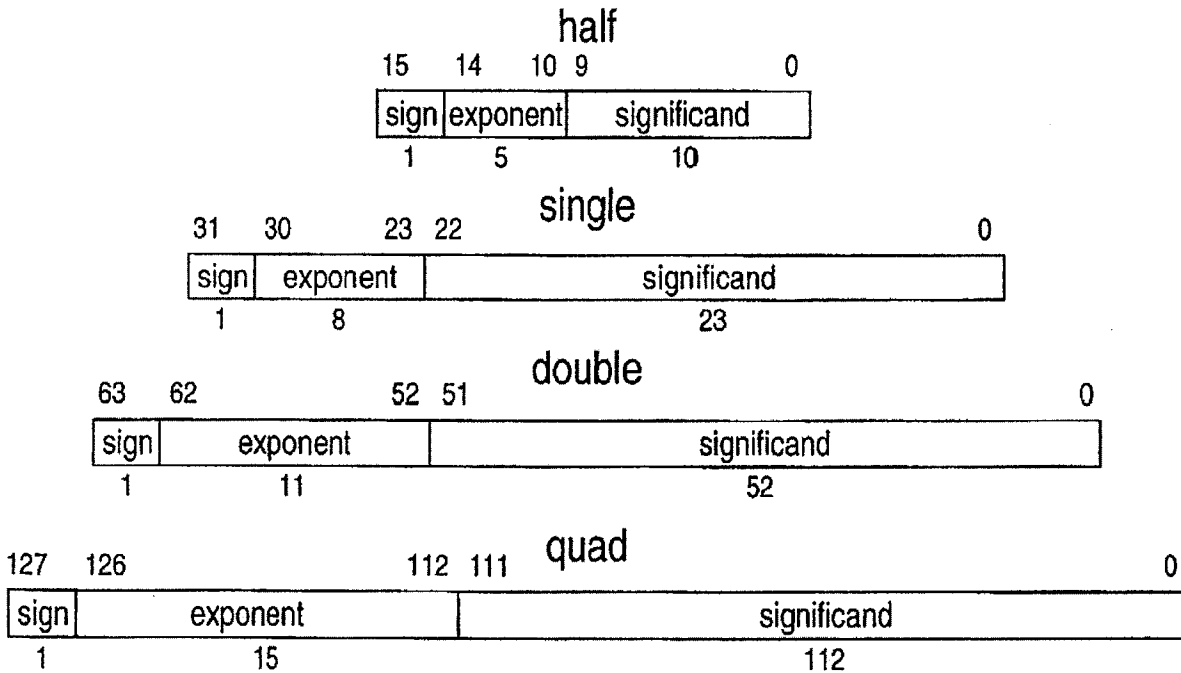
19. The Terpsichore manual describes all of the aforementioned elements of claims 9, 18, 24, 25, 31, and 32, on at least pages 19-21, 24-26, 29, 47-48, 129-131, and 150-157 (attached as Exhibit A). The Zeus manual describes all of the elements of claims 9, 18, 24, 25, 31, and 32, on at least pages 14-16, 19-21, 23-24, 55, 123-125, 128-130, and 258-260 (attached as Exhibit B).

Claims 9 and 18

20. The Terpsichore manual describes all of the aforementioned elements of claim 9 (and the substantially similar aforementioned elements of claim 18), on at least pages 19-21, 24-25, 29, 47-48, and 129-131, describing Group Floating-point instructions such as various forms of Group Floating-point Multiply instructions. Paragraphs 21-27 of this declaration discuss selected portions of those pages. Paragraphs 28-32 of this declaration discuss how the elements of claim 9 (and the substantially similar aforementioned elements of claim 18) are described by those pages. The Zeus manual describes all of the elements of claim 9 (and the substantially similar aforementioned elements of claim 18), on at least pages 14-16, 19-20, 23-24, 55, and 258-260.

Supplemental Declaration of Korbin S Van Dyke

21. The '840 patent describes various floating-point data sizes such as 16, 32, 64, and 128 bits (see '840, column 15, lines 62-65, and Fig. 9b, reproduced below). The Terpsichore manual, for example on pages 19-21, describes the various floating-point data sizes as designed to satisfy ANSI/IEEE standard 754-1985. Similarly, the Zeus manual, for example on pages 14-16, provides similar information.



22. The '840 patent provides description relating to floating-point hardware capabilities, such as in the Terpsichore manual on page 29, "operations supported in hardware are floating-point add, subtract, multiply, divide, and square root", and further on page 47, "partitioning favored for the initial implementation places all instructions that involving shifting and shuffling in one execution unit, and all instructions that involve multiplication, including fixed-point and floating-point multiply and add in another unit". Similarly, the Zeus manual, for example on pages 23-24 and 55, has similar descriptive information.

23. The Terpsichore manual describes several variations of Group Floating-point Multiply instructions, such as GF.MUL.16, GF.MUL.32, and GF.MUL.64, (among others) as described on pages 129-131, and reproduced below (excerpted and annotations added). Similarly, the Zeus manual, for example on pages 258-260, provides similar information.

Group Floating-point

These operations take two values from registers, perform floating-point arithmetic on groups of bits in the operands, and place the concatenated results in a register.

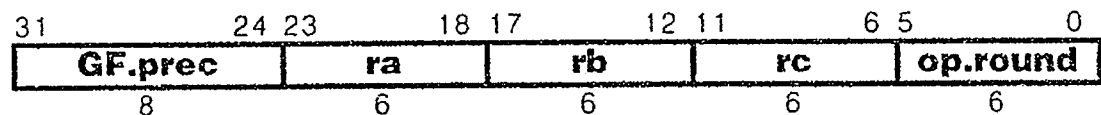
Operation codes

| | | |
|---|-------------|---|
| → | GF.MUL.16 | Group floating-point multiply half |
| | GF.MUL.16.C | Group floating-point multiply half ceiling |
| | GF.MUL.16.F | Group floating-point multiply half floor |
| | GF.MUL.16.N | Group floating-point multiply half nearest |
| | GF.MUL.16.T | Group floating-point multiply half truncate |
| | GF.MUL.16.X | Group floating-point multiply half exact |
| → | GF.MUL.32 | Group floating-point multiply single |
| | GF.MUL.32.C | Group floating-point multiply single ceiling |
| | GF.MUL.32.F | Group floating-point multiply single floor |
| | GF.MUL.32.N | Group floating-point multiply single nearest |
| | GF.MUL.32.T | Group floating-point multiply single truncate |
| | GF.MUL.32.X | Group floating-point multiply single exact |
| → | GF.MUL.64 | Group floating-point multiply double |

24. The Terpsichore manual, on page 130, describes an instruction format for the Group Floating-point instructions (including Multiply forms as well as Add and Divide forms), reproduced below. Similarly, the Zeus manual, for example on page 260, provides similar information.

Format

GF.op.prec.round rc=ra,rb

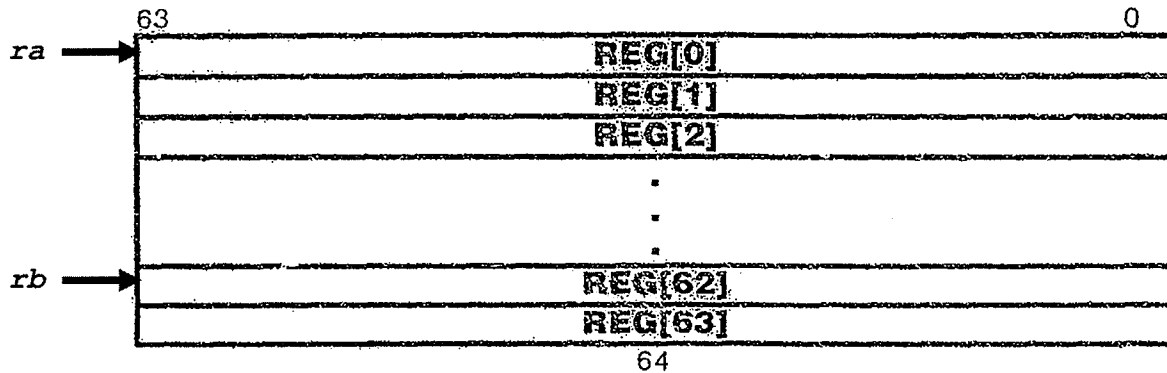


The operands of the Group Floating-point instructions (such as Multiply forms) include 'ra', 'rb', and 'rc'. As described in more detail in paragraphs 26(b)-26(c)ii of this declaration, contents of registers specified by the 'ra' and 'rb' operands are interpreted as respective collections of partitioned floating-point operands. The partitioned floating-point operands relating to 'ra' and 'rb' are pairwise multiplied together, and the results are concatenated and then stored in a register specified by the 'rc' operand.

25. The Terpsichore manual, on pages 24-25, describes registers referenced as operands of some instructions (such as Group Floating-point instructions, including Multiply forms as well as Add and Divide forms), as reproduced below (with annotations illustrating examples of an '*ra*' operand of '0' and an '*rb*' operand of '62'). Similarly the Zeus manual, for example on pages 19-20, provides similar information.

General Registers

Terpsichore user state includes 64 general registers. All are identical; there is no dedicated zero-valued register, and there are no dedicated floating-point registers.



The forgoing registers are included in register file 110 of Fig. 7 of the '840 patent.

26. The Terpsichore manual, on pages 130-131, describes a definition of various Group Floating-point instructions, including several Multiply forms. The definition is reproduced below, with annotations highlighting several elements that are discussed in the following sub-paragraphs concerning highlights of what one of ordinary skill in the art would understand from the description of the Terpsichore manual. Similarly, the Zeus manual, for example on page 260, provides similar information.

Definition

```

def GroupFloatingPoint(op,prec,round,ra,rb,rc) as
  [2] → {
    a ← RegRead(ra, 128)
    b ← RegRead(rb, 128)
    for i ← 0 to 128-prec by prec
      [4] → {
        ai ← F(prec, a, +prec-1..i)
        bi ← F(prec, b, +prec-1..i)
        if round≠NONE then
          if isSignallingNaN(ai) || isSignallingNaN(bi)
            raise FloatingPointException
          endif
          case op of
            F.DIV:
              if bi=0 then
                raise FloatingPointArithmetic
              endif
            others:
            endcase
          endif
          case op of
            GF.ADD:
              ci ← ai+bi
            [1] → GF.MUL:
              [5] → ci ← ai*bi
            GF.DIV:
              ci ← ai/bi
            endcase
          case op of
            GF.ADD, GF.MUL, GF.DIV:
              [6] → ci+prec-1..i ← PackF(prec, ci)
            endcase
          endfor
        endcase
      case round of
        X:
        N:
        T:
        F:
        C:
        NONE:
      endcase
      if rc0 then
        raise ReservedInstruction
      endif
    [7] → RegWrite(rc, 128, c)
  endcase
enddef

```

← [3]

- (a) The '*op*' field of the instruction is decoded to distinguish a Multiply form (MUL) from an Add (ADD) or Divide (DIV) form, for example, as highlighted by annotation [1].
- (b) Source operands are read from pairs of registers, as specified by the '*ra*' and '*rb*' operands, into variables '*a*' and '*b*', respectively (see annotation [2]), such as including reading REG[0] into the least-significant 64 bits of '*a*' and REG[1] into the most-significant 64 bits of '*a*', when '*ra*' is 0.

- (c) A *'for'* construct (see annotation [3]) specifies a number of evaluations of elements of the construct according to a *'prec'* operand that specifies a (floating-point) precision to interpret the operands.
- i. For example, if the *'prec'* operand is 16, then the *'for'* construct is evaluated with eight values for variable *'i'* (0, 16, 32, 48, 64, 80, 96, and 112), respectively. For another example, if the *'prec'* operand is 64, then the construct is evaluated with two values for *'i'* (0 and 64), respectively.
 - ii. Each evaluation begins by determining a partitioned floating-point value from each of the variables *'a'* and *'b'* (see annotation [4]) in accordance with the *'prec'* operand. For example, if the *'prec'* operand is 16, then for the evaluation where *'i'* is 0, a first partitioned floating-point value is determined from *'a'* from the least-significant 16 bits of *'a'*, or $a_{15..0}$ as identified by the expression $a_{i+prec-1..i}$ in the definition. Further in the evaluation where *'i'* is 0, a second floating-point value is determined from the least 16 bits of *'b'*. Continuing with the example, for the evaluation where *'i'* is 16, partitioned floating-point values are determined from the next-most-significant 16 bits of *'a'* and *'b'*, such that bits 31 to 16 determine the floating-point values. Further continuing with the example, for the evaluation where *'i'* is 112, partitioned floating-point values are determined from the most significant 16 bits of *'a'* and *'b'*. For another example of construct evaluations, if the *'prec'* operand is 64, then the *'for'* construct is evaluated with two values for *'i'* (0 and 64). Partitioned floating-point values are determined in the evaluation where *'i'* is 0 as the least-significant 64 bits of *'a'* and *'b'*, and in the evaluation where *'i'* is 64 as the most-significant 64 bits.
 - iii. The *'for'* construct processing continues by decoding a rounding mode and processing accordingly, and then decoding the *'op'* field of the instruction (see previously discussed annotation [1]).
 - iv. In the case of a Multiply form, the *'for'* construct processing continues by multiplying the partitioned floating-point values from *'a'* and *'b'* by each other (see annotation [5]). The multiplying is in accordance with floating-point multiplying.
 - v. The *'for'* construct processing completes by writing the result of the multiplies into appropriate bits of destination variable *'c'* (see annotation [6]) in accordance with the *'prec'* operand. The appropriate bit locations of *'c'* are identical to the bit locations of *'a'* and *'b'* that the partitioned floating-point values were determined from. For example, if the *'prec'* operand is 16, then for the evaluation where *'i'* is 0, the least-significant 16 bits (i.e. bits 15 to zero) of *'c'* are written, and for the evaluation where *'i'* is 16, the next-most-significant 16 bits (i.e. bits 31 to 16) of *'c'* are written. For the evaluation where *'i'* is 112, the most-significant bits (i.e. bits 127 to 112) of *'c'* are written.

Supplemental Declaration of Korbin S Van Dyke

- vi. Note that each evaluation of the '*for*' construct is independent of the other evaluations, serving to operate on different unique and non-overlapping partitioned fields of the operands. Thus each evaluation is performable in parallel with the other evaluations, sequentially with the other evaluations, or any combination thereof.

- (d) After completion of the '*for*' construct, processing completes by writing '*c*' into a pair of registers, as specified by the '*rc*' operand (see annotation [7]).

One of ordinary skill in the art would readily understand that the computation of the floating-point multiplies would occur in ALU 102 of Fig. 7.

27. Thus the Group Floating-point Multiply instructions are described in the Terpsichore manual (and also the Zeus manual) as interpreting contents of two source registers as respective pluralities of floating-point operands that are multiplied together, producing a plurality of products as results. The results are concatenated together and stored in a register specified by a third operand.

28. At least the Terpsichore manual pages 19-21, 24-25, 29, 47-48, and 129-131 describe all elements of claim 9 and substantially similar claim 18 (as amended), as described in more detail in paragraphs 29-32 of this declaration. Similarly, at least the Zeus manual pages 14-16, 19-20, 23-24, 55, and 258-260 describe all elements of claim 9 and substantially similar claim 18 (as amended).

29. The element (of claim 9) wherein the execution unit is further operable, in response to decoding a second single instruction specifying a register containing a first plurality of floating-point operands and another register containing a second plurality of floating-point operands is described by the Terpsichore manual, for example as annotated and discussed by paragraphs 21-27 of this declaration. Each of the Group Floating-point Multiply instructions is an exemplary single instruction that specifies registers containing respective pluralities of floating-point operands, via, for example, the '*ra*', and '*rb*' operands. See paragraphs 26(b) and 26(c)ii of this declaration for additional detailed discussion. As discussed in paragraph 22 of this declaration, an execution unit is clearly disclosed that is operable as claimed in this element of claim 9, as well as the other elements of claim 9 discussed in the following paragraphs.

30. The element (execution unit operable to) multiply the first plurality of floating-point operands by the second plurality of floating-point operands to produce a plurality of products is described by the Terpsichore manual, for example as annotated and discussed in paragraphs 21-27 of this declaration. Each of the Group Floating-point Multiply instructions is operable to perform a floating-point multiply on operands from registers specified by '*ra*' and '*rb*'. See paragraph 26(c)iv of this declaration for additional detailed discussion.

31. The element (execution unit operable to) provide the plurality of products to partitioned fields of a result register as a concatenated result is described by the Terpsichore manual, for example as annotated and discussed in paragraphs 21-27 of this declaration. Each of the Group Floating-point Multiply instructions is operable to produce products into fields of bits of

Supplemental Declaration of Korbin S Van Dyke

result registers via, for example, the 'rc' operand. See paragraphs 26(c)v and 26(d) of this declaration for additional detailed discussion.

32. Thus every element of claim 9 and substantially similar claim 18 (as amended) are described at least by the Terpsichore manual on pages 19-21, 24-25, 29, 47-48, and 129-131. In addition, every element of claim 9 and substantially similar claim 18 are also described at least by the Zeus manual on pages 14-16 (describing floating-point data formats), pages 19-20 (describing general registers), pages 23-24 and 55 (describing floating-point arithmetic hardware), and pages 258-260 (a definition for Ensemble Floating-point instructions, including Multiply forms).

Claims 24 and 31

33. The Terpsichore manual describes all of the aforementioned elements of claim 24 (and the substantially similar aforementioned elements of claim 31), on at least pages 24-26 and 150-157, describing Store Immediate instructions such as various forms of Store Multiplex Immediate instructions. The Zeus manual describes all of the elements of claim 24 (and the substantially similar aforementioned elements of claim 31), on at least pages 19-21, 123-125, and 128-130.

34. Claim 24 is dependent upon claim 19, and context associated with the element (of claim 24) the first and second operands is from claim 19, reproduced below (as amended):

{claim 19} A programmable processor comprising:
a virtual memory addressing unit;
an instruction path and a data path;
an external interface operable to receive data from an external
source and communicate the received data over the data path;
a cache operable to retain data communicated between the external
interface and the data path;
a register file comprising a plurality of registers coupled to the data
path; and an execution unit, coupled to the instruction and data paths, that is
operable to decode and execute instructions received from the instruction path, the
execution unit capable of performing a bitwise insert operation that operates on a
first and a second operand stored in at least one register in the register file,
wherein each bit in the second operand is individually selectable as either having
a first predetermined value or a second predetermined value, wherein for each bit
in the first operand, the bitwise insert operation inserts the bit into a
corresponding bit position in a destination value if a corresponding bit in the
second operand has the first predetermined value.

35. As is described in more detail in paragraphs 36-39 of this declaration, the Terpsichore manual, on at least pages 24-26 and 150-157, describes all of the limitations of claim 19, in addition to claim 24, with respect to several variations of Store and Store Immediate instructions, including Store Multiplex and Store Multiplex Immediate forms. The initial Van

Supplemental Declaration of Korbin S Van Dyke

Dyke declaration, in paragraphs 22-28, discusses selected portions of those pages. Similarly, at least the Zeus manual pages 19-21, 123-125, and 128-130 describe all elements of claims 19 and 24.

36. The element of (claim 19) to decode and execute instructions received from the instruction path, the execution unit capable of performing a bitwise insert operation that operates on a first and a second operand stored in at least one register in the register file is described by the Terpsichore manual, as annotated and discussed in paragraphs 22-28 of the initial Van Dyke declaration. Each of the Store Multiplex Immediate instructions is a single instruction, as evidenced at least by the dedicated operation codes “S.MUX.64.B.A.I” and “S.MUX.64.L.A.I”. Each of the Store Multiplex Immediate instructions is for performing a bitwise insert operation, since the combination of bit-wise logical-AND and bit-wise logical OR described for computing the store value results in insertion of a bit in place of another bit, e.g. “insertion” (see paragraph 27 of the initial Van Dyke declaration). The claimed at least one register corresponds to the register pair identified by the ‘rb’ operand (see paragraph 26 of the initial Van Dyke declaration). The claimed first and second operands correspond respectively to the odd- and even-numbered registers of the register pair.

37. The element (of claim 19) wherein each bit in the second operand is individually selectable as either having a first predetermined value or a second predetermined value, wherein for each bit in the first operand, the bitwise insert operation inserts the bit into a corresponding bit position in a destination value if a corresponding bit in the second operand has the first predetermined value is described by the Terpsichore manual, as annotated and discussed in paragraphs 22-28 of the initial Van Dyke declaration. As discussed in paragraph 27 of the initial Van Dyke declaration, a store value is determined by bit-wise multiplexing (e.g. “inserting”) between a first data input and a second data input, based on a control input. The second data input is from memory. The first data input is the upper 64 bits of a value identified in the Terpsichore manual as ‘m’, and the control input is the lower 64 bits of ‘m’. As discussed in paragraph 26 of the initial Van Dyke declaration, the upper 64 bits of ‘m’ are obtained from the odd-numbered register identified by ‘rb’, corresponding to the claimed first operand. Further the lower 64 bits of ‘m’ are obtained from the even-numbered register identified by ‘rb’, corresponding to the claimed second operand.

38. Therefore, according to the discussion in paragraphs 36-37 of this declaration, the element (of claim 24) wherein each of the first and second operands has a width of 64 bits the first operand corresponds to the upper 64 bits of ‘m’ that are obtained from the odd-numbered register identified by ‘rb’. The second operand corresponds to the lower 64 bits of ‘m’ that are obtained from the even-numbered register identified by ‘rb’. Thus both the first and the second operands are described as having a width of 64 bits.

39. Thus every element of claim 24 is described at least by the Terpsichore manual on pages 24-26 and 150-157. In addition, every element of claim 24 is also described at least by the Zeus manual on pages 19-21, 123-125, and 128-130.

40. Claim 31 is dependent upon claim 26, and context associated with the element (of claim 31) the first and second operands is from claim 26, reproduced below (as amended):

Supplemental Declaration of Korbin S Van Dyke

{claim 26} A device having installed therein a programmable processor, the programmable processor comprising:
a virtual memory addressing unit;
an instruction path and a data path;
an external interface operable to receive data from an external source and communicate the received data over the data path;
a cache operable to retain data communicated between the external interface and the data path;
a register file comprising a plurality of registers coupled to the data path; and an execution unit, coupled to the instruction and data paths, that is operable to decode and execute instructions received from the instruction path, the execution unit capable of performing a bitwise insert operation that operates on a first and a second operand stored in at least one register in the register file, wherein each bit in the second operand is individually selectable as either having a first predetermined value or a second predetermined value, wherein for each bit in the first operand, the bitwise insert operation inserts the bit into a corresponding bit position in a destination value if a corresponding bit in the second operand has the first predetermined value.

41. Claim 26 (a device claim) is substantially similar to claim 19 (a programmable processor claim), and as previously discussed, claim 31 is substantially similar to claim 24. Therefore paragraphs 33-39 in this declaration concerning claim 24 and parent claim 19 are also applicable to claim 31 and parent claim 26. Thus every element of claim 31 is described at least by the Terpsichore manual on pages 24-26 and 150-157. In addition, every element of claim 31 is also described at least by the Zeus manual on pages 19-21, 123-125, and 128-130.

Claims 25 and 32

42. The Terpsichore manual describes all of the aforementioned elements of claim 25 (and the substantially similar aforementioned elements of claim 32), on at least pages 19-21, 24-25, 29, 47-48, and 129-131, describing Group Floating-point instructions such as various forms of Group Floating-point Add, Divide, and Multiply instructions. Paragraphs 43-50 of this declaration discuss selected portions of those pages. Paragraphs 51-52 of this declaration describe how the elements of claim 25 (and the substantially similar aforementioned elements of claim 32) are described by those pages. The Zeus manual describes all of the elements of claim 25 (and the substantially similar aforementioned elements of claim 32), on at least pages 14-16, 19-20, 23-24, 55, and 258-260.

43. As discussed in more detail in paragraph 21 of this declaration, the '840 and the '599 patents describe various floating-point data sizes.

44. As discussed in more detail in paragraph 22 of this declaration, the '840 and the '599 patents describe various floating-point hardware capabilities.

45. The Terpsichore manual describes several types and variations of Group Floating-point instructions, such as Group Floating-point Add (e.g. GF.ADD.64), Divide (e.g. GF.DIV.32), and Multiply (e.g. GF.MUL.16) instructions, as described on pages 129-131, and reproduced below (excerpted and annotations added). Similarly, the Zeus manual, for example on pages 258-260, provides similar information.

Group Floating-point

These operations take two values from registers, perform floating-point arithmetic on groups of bits in the operands, and place the concatenated results in a register.

Operation codes

| | |
|--------------|---|
| → GF.ADD.64 | Group floating-point add double |
| GF.ADD.64 .C | Group floating-point add double ceiling |
| GF.ADD.64 .F | Group floating-point add double floor |
| GF.ADD.64 .N | Group floating-point add double nearest |
| GF.ADD.64 .T | Group floating-point add double truncate |
| GF.ADD.64 .X | Group floating-point add double exact |
| GF.DIV.16 | Group floating-point divide half |
| GF.DIV.16.C | Group floating-point divide half ceiling |
| GF.DIV.16.F | Group floating-point divide half floor |
| GF.DIV.16.N | Group floating-point divide half nearest |
| GF.DIV.16.T | Group floating-point divide half truncate |
| GF.DIV.16.X | Group floating-point divide half exact |
| → GF.DIV.32 | Group floating-point divide single |
| GF.DIV.32.C | Group floating-point divide single ceiling |
| GF.DIV.32.F | Group floating-point divide single floor |
| GF.DIV.32.N | Group floating-point divide single nearest |
| GF.DIV.32.T | Group floating-point divide single truncate |
| GF.DIV.32.X | Group floating-point divide single exact |
| GF.DIV.64 | Group floating-point divide double |
| GF.DIV.64.C | Group floating-point divide double ceiling |
| GF.DIV.64.F | Group floating-point divide double floor |
| GF.DIV.64.N | Group floating-point divide double nearest |
| GF.DIV.64.T | Group floating-point divide double truncate |
| GF.DIV.64.X | Group floating-point divide double exact |
| → GF.MUL.16 | Group floating-point multiply half |
| GF.MUL.16.C | Group floating-point multiply half ceiling |
| GF.MUL.16.F | Group floating-point multiply half floor |

46. As discussed in paragraph 24 of this declaration, the '840 and the '599 patents describe an instruction format for the Group Floating-point instructions (including Add, Divide, and Multiply forms).

Supplemental Declaration of Korbin S Van Dyke

47. As discussed in paragraph 25 of this declaration, the '840 and the '599 patents describe registers referenced as operands of some instructions (such as Group Floating-point Add, Divide, and Multiply instructions).

48. As discussed in paragraph 26 of this declaration, the '840 patent and the '599 patent describe definitions of various Group Floating-point instructions, including several Multiply forms. In addition, Add and Divide forms of Group Floating-point instructions are described by the '840 patent and the '599 patent, as evidenced by the following excerpt from page 131 of the Terpsichore manual:

```
case op of
  GF.ADD:
    ci ← ai+bi
  GF.MUL:
    ci ← ai*bi
  GF.DIV.:
    ci ← ai/bi
endcase
```

Similarly, the Zeus manual, for example on page 260, provides similar information.

49. The discussion of Multiply forms of Group Floating-point instructions in subparagraphs 26(a)-26(d) of this declaration is generally applicable to Add and Divide forms, as well. Rather than multiplying the partitioned floating-point values (as discussed in subparagraph 26(c)iv of this declaration), the floating-point values are added (Add form) or divided (Divide form).

50. Thus the Group Floating-point instructions of Add, Divide, and Multiply forms are described in the Terpsichore manual (and also the Zeus manual) as interpreting contents of two source registers as respective pluralities of floating-point operands that are added, divided, or multiplied together, respectively, producing a plurality of floating-point results that are concatenated together and stored in a register specified by a third operand.

51. The element (of claim 25) wherein the execution unit is further capable of executing a plurality of different group floating-point arithmetic operations that arithmetically operate on multiple floating-point operands stored in partitioned fields of registers in the register file to produce a catenated result that is returned to a register in the register file, wherein the catenated result comprises a plurality of individual floating-point results is described by the Terpsichore manual, for example as annotated and discussed in paragraphs 43-50 of this declaration. The Add, Divide, and Multiply forms of Group Floating-point instructions are exemplary instructions that embody arithmetic operations that operate on multiple floating-point operands, for example as specified by the 'ra', and 'rb' operands. Results of the arithmetic operations are returned, for example, to result registers specified by the 'rc' operand. As discussed in paragraph 22 of this declaration, an execution unit is clearly disclosed that is operable as claimed in claim 25.

52. Thus every element of claim 25 and substantially similar claim 32 (as amended) are described at least by the Terpsichore manual on pages 19-21, 24-25, 29, 47-48, and 129-131.

Supplemental Declaration of Korbin S Van Dyke

In addition, every element of claim 25 and substantially similar claim 32 are also described at least by the Zeus manual on pages 14-16 (describing floating-point data formats), pages 19-20 (describing general registers), pages 23-24 and 55 (describing floating-point arithmetic hardware), and pages 258-260 (a definition for Ensemble Floating-point instructions, including Add, Divide, and Multiply forms).

Summary and Closing:

53. The '840 patent, including the Terpsichore manual, provides sufficient information in sufficient detail describing the claimed invention (as amended) of the 10/757,516 patent application, that one of ordinary skill in the art would reasonably conclude that the inventors had possession of the claimed invention at the time of filing the '840 patent. Further, the '840 patent, including the Terpsichore manual, provides sufficient information regarding the subject matter of the claimed invention (as amended) of the 10/757,516 patent application to enable one of ordinary skill in the pertinent art to make and use the claimed invention without undue experimentation. In addition, the '599 patent, including the Zeus manual, provides sufficient information in sufficient detail describing the claimed invention (as amended) of the 10/757,516 patent application, that one of ordinary skill in the art would reasonably conclude that the inventors had possession of the claimed invention at the time of filing the '599 patent. Further, the '599 patent, including the Zeus manual, provides sufficient information regarding the subject matter of the claimed invention (as amended) of the 10/757,516 patent application to enable one of ordinary skill in the pertinent art to make and use the claimed invention without undue experimentation.

54. Therefore, I believe that each of the '840 patent, including the Terpsichore manual, and the '599 patent, including the Zeus manual, provide adequate written description and enablement as required by 35 USC § 112 for the limitations of claims 9, 18, 24, 25, 31, and 32 (as amended) of the 10/757,516 patent application, as discussed in paragraph 3 of this declaration.

55. I have had no communication with any of the inventors of the 10/757,516 patent application (Craig Hansen and John Moussouris) relating to any material in this declaration.

56. I have been hired as a consultant in connection with procedures before the United States Patent and Trademark Office (USPTO) regarding patents and patent applications assigned to Microunity Systems Engineering, Inc., including the media processor patent application. I am being compensated for my services at the rate of \$325/hour. Other than acting as a consultant in connection with procedures before the USPTO, I have no interest or connection with Microunity Systems Engineering, Inc.

57. During my evaluation of the media processor patent application, I have been impressed by the thoroughness and overall high-quality of the Terpsichore and Zeus manuals. The manuals provide clear and unambiguous descriptions of media processing systems and are thorough and well-written. The manuals provide comprehensive descriptions of instructions in complete architectural detail. The information in the manuals would have been readily understood and easily accessible to software engineers coding the media processing systems, and hardware engineers implementing microprocessors for use in the media processing systems, and

Supplemental Declaration of Korbin S Van Dyke

that is exactly what architecture reference manuals should be. This is not surprising, since the '840 patent and the '599 patent each include an architecture manual that is intended to enable hardware engineers to do exactly that – design, build, and implement a media processor that would include circuitry for the claim limitations set forth in paragraph 3 of this declaration, as described in the Terpsichore and the Zeus architecture manuals.

Supplemental Declaration of Korbin S Van Dyke

58. I hereby declare that all statements made herein are of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issuing therefrom.

Date: 4/24/2007

Korbin Van Dyke: K. Van Dyke

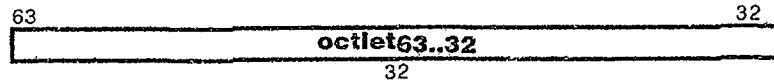
Address: 3343 Little Valley Rd.
Sunnyvale, CA 94586

Terpsichore System Architecture

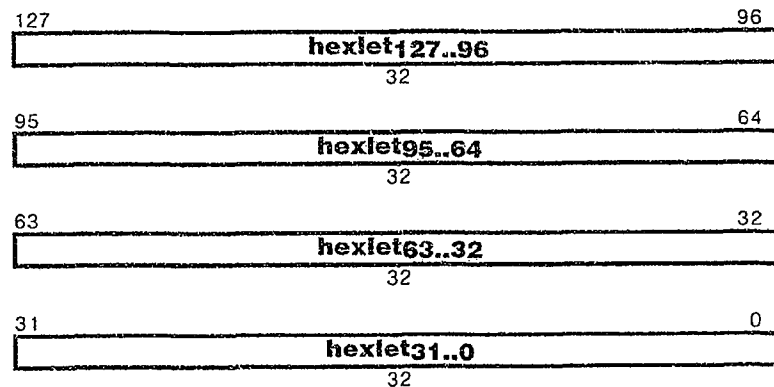
Wed, Aug 2, 1995

Octlet

An octlet is the catenation of 64 bits, and is the concatenation of eight bytes:

Hexlet

A hexlet is the catenation of 128 bits, and is the concatenation of sixteen bytes:

Address

Terpsichore addresses are octlet quantities.

Floating-point Data

Terpsichore's floating-point formats are designed to satisfy ANSI/IEEE standard 754-1985: Binary Floating-point Arithmetic. Standard 754 leaves certain aspects to the discretion of the implementor:

Terpsichore adds additional half-precision and quad-precision formats to standard 754's single-precision and double-precision formats. Terpsichore's double-precision satisfies standard 754's precision requirements for a single-extended format, and Terpsichore's quad-precision satisfies standard 754's precision requirements for a double-extended format.

Quiet NaN values are denoted by any sign bit value, an exponent field of all one bits, and a non-zero significand with the most significant bit cleared. Quiet NaN

Terpsichore System Architecture

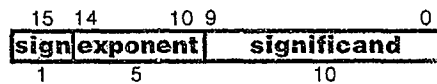
Wed, Aug 2, 1995

values generated by default exception handling of standard operations have a zero sign bit, an exponent field of all one bits, and a significand field with the most significant bit cleared, the next-most significant bit set, and all other bits cleared.

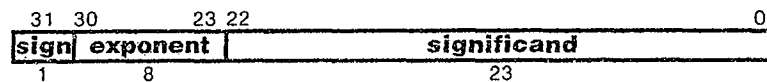
Signaling NaN values are denoted by any sign bit value, an exponent field of all one bits, and a non-zero significand with the most significant bit set.

Half-precision Floating-point

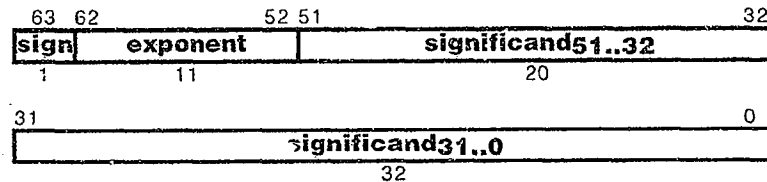
Terpsichore half precision uses a format similar to standard 754's requirements, reduced to a 16-bit overall format. The format contains sufficient precision and exponent range to hold a 12-bit signed integer.

Single-precision Floating-point

Terpsichore single precision satisfies standard 754's requirements for "single."

Double-precision Floating-point

Terpsichore double precision satisfies standard 754's requirements for "double."

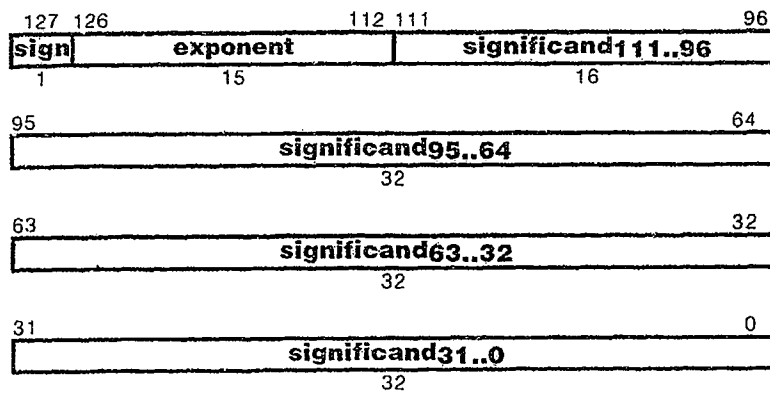


Terpsichore System Architecture

Wed, Aug 2, 1995

Quad-precision Floating-point

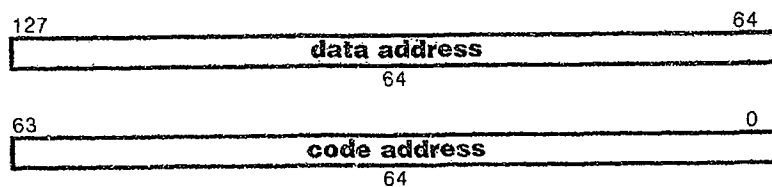
Terpsichore quad precision satisfies standard 754's requirements for "double extended," but has additional significand precision to use 128 bits.



Terpsichore System Architecture

Wed, Aug 2, 1995

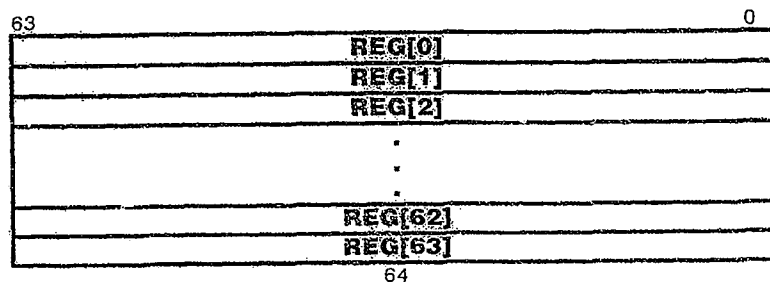
The gateway contains two data items within its structure, a code address and a data address:

User state

The user state consists of hardware data structures that are accessible to all conventional compiled code. The Terpsichore user state is designed to be as regular as possible, and consists only of the general registers, the program counter, and virtual memory. There are no specialized registers for condition codes, operating modes, rounding modes, integer multiply/divide, or floating-point values.

General Registers

Terpsichore user state includes 64 general registers. All are identical; there is no dedicated zero-valued register, and there are no dedicated floating-point registers.

Definition

```

def val ← RegRead(rn, size)
  case size of
    64:
      val ← REG[rn]
    128:
      if rn0 then
        raise ReservedInstruction
      endif
      val ← REG[rn+1] || REG[rn]
  endcase
enddef

```

Terpsichore System Architecture

Wed, Aug 2, 1995

```

def RegWrite(rn, size, val)
  case size of
    64:
      REG[rn] ← val63..0
    128:
      if rn0 then
        raise ReservedInstruction
      endif
      REG[rn+1] ← val127..64
      REG[rn] ← val63..0
  endcase
enddef

```

Program Counter

The program counter contains the address of the currently executing instruction. This register is implicitly manipulated by branch instructions, and read by branch instructions that save a return address in a general register.

Privilege Level

The privilege level register contains the privilege level of the currently executing instruction. This register is implicitly manipulated by branch gateway and branch down instructions, and read by branch gateway instructions that save a return address in a general register.

Program Counter and Privilege Level

The program counter and privilege level may be packed into a single octet. This combined data structure is saved by the Branch Gateway instruction and restored by the Branch Down instruction.

System state

The system state consists of the facilities not normally used by conventional compiled code. These facilities provide mechanisms to execute such code in a fully virtual environment. All system state is memory mapped, so that it can be manipulated by compiled code.

Fixed-point

Terpsichore provides load and store instructions to move data between memory and the registers, branch instructions to compare the contents of registers and to transfer control from one code address to another, and arithmetic operations to perform computation on the contents of registers, returning the result to registers.

Load and Store

The load and store instructions move data between memory and the registers. When loading data from memory into a register, values are zero-extended or sign-extended to fill the register. When storing data from a register into memory, values are truncated on the left to fit the specified memory region.

Load and store instructions that specify a memory region of more than one byte may use either little-endian or big-endian byte ordering; the size and ordering are explicitly specified in the instruction. Regions larger than one byte may be either aligned to addresses that are an even multiple of the size of the region, or of unspecified alignment: alignment checking is also explicitly specified in the instruction.

The load and store instructions are used for fixed-point data as well as floating-point and digital signal processing data; Terpsichore has a single bank of registers for all data types.

Swap instructions provide multithread and multiprocessor synchronization, using indivisible operations: add-and-swap, compare-and-swap, and multiplex-and-swap. A store-multiplex operation provides the ability to indivisibly write to a portion of an octlet. These instructions always operate on aligned octlet data, using either little-endian or big-endian byte ordering.

Branch Conditionally

The fixed-point compare-and-branch instructions provide all arithmetic tests for equality and inequality of signed and unsigned fixed-point values. Tests are performed either between two operands contained in general registers, or on the bitwise and of two operands. Depending on the result of the compare, either a branch is taken, or not taken. A taken branch causes an immediate transfer of the program counter to the target of the branch, specified by a 12-bit signed offset from the location of the branch instruction. A non-taken branch causes no transfer; execution continues with the following instruction.

Branch Unconditionally

Other branch instructions provide for unconditional transfer of control to addresses too distant to be reached by a 12-bit offset, and to transfer to a target while placing the location following the branch into a register. The branch through gateway instruction provides a secure means to access code at a higher privilege level, in a form similar to a normal procedure call.

Terpsichore System Architecture

Wed, Aug 2, 1995

Arithmetic Operations

The operations supported in hardware are floating-point add, subtract, multiply, divide, and square root. Other operations required by the ANSI/IEEE floating-point standard are provided by software libraries.

The operations explicitly specify the precision of the operation, and round the result to the specified precision at the conclusion of each operation.

A single instruction provides a floating-point multiply with the result fed into a floating-point add. The result is computed as if the multiply is performed to infinite precision, added as if in infinite precision, then rounded. This operation is a particularly good match to the needs of vector linear algebra routines.

Rounding

Rounding is specified within the instructions explicitly, to avoid maintaining explicit state for a rounding mode.

Exceptions

All the mandated floating-point exception conditions cause a trap when they occur; maintenance of sticky and other status bits may be performed using software routines. Because the floating-point inexact exception may be very frequent, this exception only occurs when specified in the instruction explicitly. Arithmetic operations may also specify that all exceptions are to be handled by default, generating special results instead of traps.

Digital Signal Processing

The Terpsichore processor provides a set of operations that maintain the fullest possible use of 64- and 128-bit data paths when operating on lower-precision fixed-point or floating-point vector values. These operations are useful for several application areas, including digital signal processing, image processing, and synthetic graphics. The basic goal of these operations is to accelerate the performance of algorithms that exhibit the following characteristics:

Low-precision arithmetic

The operands and intermediate results are fixed-point values represented in no greater than 64 bit precision. For floating-point arithmetic, operands and intermediate results are of 16, 32, or 64 bit precision.

The use of fixed-point arithmetic permits various forms of operation reordering that are not permitted in floating-point arithmetic. Specifically, commutativity and associativity, and distribution identities can be used to reorder operations. Compilers can evaluate operations to determine what intermediate precision is required to get the specified arithmetic result.

Terpsichore System Architecture

Wed, Aug 2, 1995

branch predicts that a future execution of the same branch will be taken. More elaborate prediction may cache the source and target addresses of multiple branches, both conditional and unconditional, and both forward and reverse.

The hardware prediction mechanism is tuned for optimizing conditional branches that close loops or express frequent alternatives, and will generally require substantially more cycles when executing conditional branches whose outcome is not predominately taken or not-taken. For such cases of unpredictable conditional results, the use of code which avoids conditional branches in favor of the use of set on compare and multiplex instructions may result in greater performance.

Where the above technique may not be applicable, a Euterpe pipeline may ensure that conditional branches which have a small positive offset be handled as if the branch is always predicted to be not taken, with the recovery of a misprediction causing cancellation of the instructions which have already been issued but not completed which would be skipped over by the taken conditional branch. This "conditional-skip" optimization is performed by the Euterpe implementation and requires no specific architectural feature to access or implement.

A Euterpe pipeline may also perform "branch-return" optimization, in which a branch-and-link instruction saves a branch target address which is used to predict the target of the next branch-register instruction. This optimization may be implemented with a depth of one (only one return address kept), or as a stack of finite depth, where a branch and link pushes onto the stack, and a branch-register pops from the stack. This optimization can eliminate the misprediction cost of simple procedure calls, as the calling branch is susceptible to hardware prediction, and the returning branch is predictable by the branch-return optimization. Like the conditional-skip optimization described above, this feature is performed by the Euterpe implementation and requires no specific architectural feature to access or implement.

Additional Load and Execute Resources

Studies of the dynamic distribution of Euterpe instructions on various benchmark suites indicate that the most frequently-issued instruction classes are load instructions and execute instructions. In a high-performance Euterpe implementation, it is advantageous to consider execution pipelines in which the ability to target the machine resources toward issuing load and execute instructions is increased.

One of the means to increase the ability to issue execute-class instructions is to provide the means to issue two execute instructions in a single-issue string. The execution unit actually requires several distinct resources, so by partitioning these resources, the issue capability can be increased without increasing the number of functional units, other than the increased register file read and write ports. The partitioning favored for the initial implementation places all instructions that involve shifting and shuffling in one execution unit, and all instructions that involve multiplication, including fixed-point and floating-point multiply and add in another unit. Resources used for implementing add, subtract, and bitwise logical operations may be duplicated, being modest in size compared to the shift and

multiply units, or shared between the two units, as the operations have low-enough latency that two operations might be pipelined within a single issue cycle. These instructions must generally be independent, except perhaps that two simple add, subtract, or bitwise logical may be performed dependently, if the resources for executing simple instructions are shared between the execution units.

One of the means to increase the ability to issue load-class instructions is to provide the means to issue two load instructions in a single-issue string. This would generally increase the resources required of the data fetch unit and the data cache, but a compensating solution is to steal the resources for the store instruction to execute the second load instruction. Thus, a single-issue string can then contain either two load instructions, or one load instruction and one store instruction, which uses the same register read ports and address computation resources as the basic 5-instruction string. This capability also may be employed to provide support for unaligned load and store instructions, where a single-issue string may contain as an alternative a single unaligned load or store instruction which uses the resources of the two load-class units in concert to accomplish the unaligned memory operation.

Result Forwarding

When temporally adjacent instructions are executed by separate resources, the results of the first instruction must generally be forwarded directly to the resource used to execute the second instruction, where the result replaces a value which may have been fetched from a register file. Such forwarding paths use significant resources. A Terpsichore implementation must generally provide forwarding resources so that dependencies from earlier instructions within a string are immediately forwarded to later instructions, except between a first and second execution instruction as described above. In addition, when forwarding results from the execution units back to the data fetch unit, additional delay may be incurred.

Group Floating-point

These operations take two values from registers, perform floating-point arithmetic on groups of bits in the operands, and place the concatenated results in a register.

Operation codes

| | |
|-------------|---|
| GF.ADD.13 | Group floating-point add half |
| GF.ADD.16.C | Group floating-point add half ceiling |
| GF.ADD.16.F | Group floating-point add half floor |
| GF.ADD.16.N | Group floating-point add half nearest |
| GF.ADD.16.T | Group floating-point add half truncate |
| GF.ADD.16.X | Group floating-point add half exact |
| GF.ADD.32 | Group floating-point add single |
| GF.ADD.32.C | Group floating-point add single ceiling |
| GF.ADD.32.F | Group floating-point add single floor |
| GF.ADD.32.N | Group floating-point add single nearest |
| GF.ADD.32.T | Group floating-point add single truncate |
| GF.ADD.32.X | Group floating-point add single exact |
| GF.ADD.64 | Group floating-point add double |
| GF.ADD.64.C | Group floating-point add double ceiling |
| GF.ADD.64.F | Group floating-point add double floor |
| GF.ADD.64.N | Group floating-point add double nearest |
| GF.ADD.64.T | Group floating-point add double truncate |
| GF.ADD.64.X | Group floating-point add double exact |
| GF.DIV.16 | Group floating-point divide half |
| GF.DIV.16.C | Group floating-point divide half ceiling |
| GF.DIV.16.F | Group floating-point divide half floor |
| GF.DIV.16.N | Group floating-point divide half nearest |
| GF.DIV.16.T | Group floating-point divide half truncate |
| GF.DIV.16.X | Group floating-point divide half exact |
| GF.DIV.32 | Group floating-point divide single |
| GF.DIV.32.C | Group floating-point divide single ceiling |
| GF.DIV.32.F | Group floating-point divide single floor |
| GF.DIV.32.N | Group floating-point divide single nearest |
| GF.DIV.32.T | Group floating-point divide single truncate |
| GF.DIV.32.X | Group floating-point divide single exact |
| GF.DIV.64 | Group floating-point divide double |
| GF.DIV.64.C | Group floating-point divide double ceiling |
| GF.DIV.64.F | Group floating-point divide double floor |
| GF.DIV.64.N | Group floating-point divide double nearest |
| GF.DIV.64.T | Group floating-point divide double truncate |
| GF.DIV.64.X | Group floating-point divide double exact |
| GF.MUL.16 | Group floating-point multiply half |
| GF.MUL.16.C | Group floating-point multiply half ceiling |
| GF.MUL.16.F | Group floating-point multiply half floor |

Terpsichore System Architecture

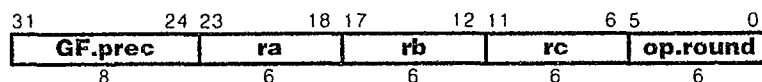
Wed, Aug 2, 1995

| | |
|-------------|---|
| GF.MUL.16.N | Group floating-point multiply half nearest |
| GF.MUL.16.T | Group floating-point multiply half truncate |
| GF.MUL.16.X | Group floating-point multiply half exact |
| GF.MUL.32 | Group floating-point multiply single |
| GF.MUL.32.C | Group floating-point multiply single ceiling |
| GF.MUL.32.F | Group floating-point multiply single floor |
| GF.MUL.32.N | Group floating-point multiply single nearest |
| GF.MUL.32.T | Group floating-point multiply single truncate |
| GF.MUL.32.X | Group floating-point multiply single exact |
| GF.MUL.64 | Group floating-point multiply double |
| GF.MUL.64.C | Group floating-point multiply double ceiling |
| GF.MUL.64.F | Group floating-point multiply double floor |
| GF.MUL.64.N | Group floating-point multiply double nearest |
| GF.MUL.64.T | Group floating-point multiply double truncate |
| GF.MUL.64.X | Group floating-point multiply double exact |

| | op | prec | | | | round/trap |
|----------|-----|------|----|----|-----|---------------|
| add | ADD | 16 | 32 | 64 | 128 | NONE CF N T X |
| divide | DIV | 16 | 32 | 64 | 128 | NONE CF N T X |
| multiply | MUL | 16 | 32 | 64 | 128 | NONE CF N T X |

Format

GF.op.prec.round rc=ra,rb

Description

The contents of registers ra and rb are combined using the specified floating-point operation. The result is placed in register rc. The operation is rounded using the specified rounding option or using round-to-nearest if not specified. If a rounding option is specified, the operation raises a floating-point exception if a floating-point invalid operation, divide by zero, overflow, or underflow occurs, or when specified, if the result is inexact. If a rounding option is not specified, floating-point exceptions are not raised, and are handled according to the default rules of IEEE 754.

Definition

```

def GroupFloatingPoint(op,prec,round,ra,rb,rc) as
  a ← RegRead(ra, 128)
  b ← RegRead(rb, 128)
  for i ← 0 to 128-prec by prec
    ai ← F(prec,a,prec-1,i)
    bi ← F(prec,b,prec-1,i)
    if round≠NONE then

```

Terpsichore System Architecture

Wed, Aug 2, 1995

```

        if isSignallingNaN(ai) | isSignallingNaN(bi)
            raise FloatingPointException
        endif
        case op of
            F.DIV:
                if bi=0 then
                    raise FloatingPointArithmetic
                endif
                others:
                    endcase
            endcase
        endif
        case op of
            GF.ADD:
                ci ← ai+bi
            GF.MUL:
                ci ← ai*bi
            GF.DIV.:
                ci ← ai/bi
        endcase
        case op of
            GF.ADD, GF.MUL, GF.DIV.:
                Ci+prec-1..i ← PackF(prec, ci)
        endcase
    endfor
endcase
case round of
    X:
    N:
    T:
    F:
    C:
    NONE:
endcase
if rc0 then
    raise ReservedInstruction
endif
RegWrite(rc, 128, c)
endcase
enddef

```

Exceptions

Reserved instruction
Floating-point arithmetic



Store

These operations add the contents of two registers to produce a virtual address, and store the contents of a register into memory.

Operation codes

| | |
|-------------------|---|
| S.8 ⁴⁶ | Store byte |
| S.16.B | Store double big-endian |
| S.16.B.A | Store double big-endian aligned |
| S.16.L | Store double little-endian |
| S.16.L.A | Store double little-endian aligned |
| S.32.B | Store quadlet big-endian |
| S.32.B.A | Store quadlet big-endian aligned |
| S.32.L | Store quadlet little-endian |
| S.32.L.A | Store quadlet little-endian aligned |
| S.64.B | Store octlet big-endian |
| S.64.B.A | Store octlet big-endian aligned |
| S.64.L | Store octlet little-endian |
| S.64.L.A | Store octlet little-endian aligned |
| S.128.B | Store hexlet big-endian |
| S.128.B.A | Store hexlet big-endian aligned |
| S.128.L | Store hexlet little-endian |
| S.128.L.A | Store hexlet little-endian aligned |
| S.AAS.64.B.A | Store add-and-swap octlet big-endian aligned |
| S.AAS.64.L.A | Store add-and-swap octlet little-endian aligned |
| S.CAS.64.B.A | Store compare-and-swap octlet big-endian aligned |
| S.CAS.64.L.A | Store compare-and-swap octlet little-endian aligned |
| S.MAS.64.B.A | Store multiplex-and-swap octlet big-endian aligned |
| S.MAS.64.L.A | Store multiplex-and-swap octlet little-endian aligned |
| S.MUX.64.B.A | Store multiplex octlet big-endian aligned |
| S.MUX.64.L.A | Store multiplex octlet little-endian aligned |

| size | ordering | alignment |
|-----------------------|----------|-----------|
| 8 | | |
| 16 32 64 128 | L B | |
| 16 32 64 128 | L B | A |

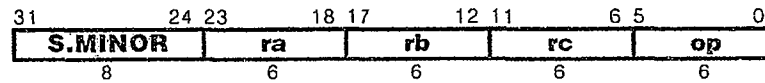
⁴⁶S.8 need not specify byte ordering, nor need it specify alignment checking, as it stores a single byte.

Terpsichore System Architecture

Wed, Aug 2, 1995

Format

op ra,rb,rc

Description

A virtual address is computed from the sum of the contents of register ra and register rb. The contents of register rc, treated as the size specified, is stored in memory using the specified byte order.

If alignment is specified, the computed virtual address must be aligned, that is, it must be an exact multiple of the size expressed in bytes. If the address is not aligned an "access disallowed by virtual address" exception occurs.

Definition

def Store(op,ra,rb,rc) as

case op of

S8,

S16L, S16LA, S16B, S16BA,

S32L, S32LA, S32B, S32BA,

S64L, S64LA, S64B, S64BA,

S128L, S128LA, S128B, S128BA:

function ← NONE

SAAS64BA, SAAS64LA:

function ← AAS

SCAS64BA, SCAS64LA:

function ← CAS

SMAS64BA, SMAS64LA:

function ← MAS

SMUX64BA, SMUX64LA:

function ← MUX

endcase

case op of

S8:

size ← 8

S16L, S16LA, S16B, S16BA:

size ← 16

S32L, S32LA, S32B, S32BA:

size ← 32

S64L, S64LA, S64B, S64BA,

SAAS64BA, SAAS64LA:

size ← 64

SCAS64BA, SCAS64LA, SMAS64BA, SMAS64LA, SMUX64BA, SMUX64LA:

size ← 64

S128L, S128LA, S128B, S128BA:

size ← 128

endcase

case op of

S8,

S16L, S16LA, S16B, S16BA,

S32L, S32LA, S32B, S32BA,

Terpsichore System Architecture

Wed, Aug 2, 1995

```

S64L, S64LA, S64B, S64BA,
SAAS64BA, SAAS64LA:
    rsize ← 64
SCAS64BA, SCAS64LA, SMAS64BA, SMAS64LA, SMUX64BA, SMUX64LA:
    rsize ← 128
S128L, S128LA, S128B, S128BA:
    rsize ← 128
endcase
case op of
    S8:
        align ← undefined
        S16L, S32L, S64L, S128L,
        S16B, S32B, S64B, S128B:
            align ← false
        S16LA, S32LA, S64LA, S128LA,
        S16BA, S32BA, S64BA, S128BA,
        SAAS64BA, SAAS64LA, SCAS64BA, SCAS64LA,
        SMAS64BA, SMAS64LA, SMUX64BA, SMUX64LA:
            align ← true
    endcase
case op of
    S8:
        order ← undefined
        S16L, S32L, S64L, S128L,
        S16LA, S32LA, S64LA, S128LA,
        SAAS64BA, SCAS64LA, SMAS64BA, SMUX64LA:
            order ← L
        S16B, S32B, S64B, S128B,
        S16BA, S32BA, S64BA, S128BA,
        SAAS64BA, SCAS64BA, SMAS64BA, SMUX64BA:
            order ← B
    endcase
a ← RegRead(ra, 64)
b ← RegRead(rb, 64)
VirtAddr ← a + b
if align then
    if (VirtAddr and ((size/8)-1)) ≠ 0 then
        raise AccessDisallowedByVirtualAddress
    endif
endif
m ← RegRead(rc, rsize)
case function of
    NONE:
        StoreMemory(VirtAddr, size, order, msize-1..0)
    AAS:
        c ← LoadMemory(VirtAddr, size, order)
        StoreMemory(VirtAddr, size, order, m63..0+c)
        RegWrite(rc, 64, c)
    CAS:
        c ← LoadMemory(VirtAddr, size, order)
        if (c = m63..0) then
            StoreMemory(VirtAddr, size, order, m127..64)
        endif
        RegWrite(rc, 64, c)
    MAS:
        c ← LoadMemory(VirtAddr, size, order)
        n ← (m127..64 & m63..0) | (c & ~m63..0)
        StoreMemory(VirtAddr, size, order, n)

```

Terpsichore System Architecture

Wed, Aug 2, 1995

```

      RegWrite(rc, 64, c)
MUX:
  c ← LoadMemory(VirtAddr, size, order)
  n ← (m127..64 & m63..0) | (c & ~m63..0)
  StoreMemory(VirtAddr, size, order, n)
endcase
enddef

```

Exceptions

Reserved instruction
 Access disallowed by virtual address
 Access disallowed by tag
 Access disallowed by global TLB
 Access disallowed by local TLB
 Access detail required by tag
 Access detail required by local TLB
 Access detail required by global TLB
 Cache coherence intervention required by tag
 Cache coherence intervention required by local TLB
 Cache coherence intervention required by global TLB
 Local TLB miss
 Global TLB miss

Store Immediate

These operations add the contents of a register to a sign-extended immediate value to produce a virtual address, and store the contents of a register into memory.

Operation codes

| | |
|---------------------|---|
| S.8.I ⁴⁷ | Store byte immediate |
| S.16.B.A.I | Store double big-endian aligned immediate |
| S.16.B.I | Store double big-endian immediate |
| S.16.L.A.I | Store double little-endian aligned immediate |
| S.16.L.I | Store double little-endian immediate |
| S.32.B.A.I | Store quadlet big-endian aligned immediate |
| S.32.B.I | Store quadlet big-endian immediate |
| S.32.L.A.I | Store quadlet little-endian aligned immediate |
| S.32.L.I | Store quadlet little-endian immediate |
| S.64.B.A.I | Store octlet big-endian aligned immediate |
| S.64.B.I | Store octlet big-endian immediate |
| S.64.L.A.I | Store octlet little-endian aligned immediate |
| S.64.L.I | Store octlet little-endian immediate |
| S.128.B.A.I | Store hexlet big-endian aligned immediate |
| S.128.B.I | Store hexlet big-endian immediate |
| S.128.L.A.I | Store hexlet little-endian aligned immediate |
| S.128.L.I | Store hexlet little-endian immediate |
| S.AAS.64.B.A.I | Store add-and-swap octlet big-endian aligned immediate |
| S.AAS.64.L.A.I | Store add-and-swap octlet little-endian aligned immediate |
| S.CAS.64.B.A.I | Store compare-and-swap octlet big-endian aligned immediate |
| S.CAS.64.L.A.I | Store compare-and-swap octlet little-endian aligned immediate |
| S.MAS.64.B.A.I | Store multiplex-and-swap octlet big-endian aligned immediate |
| S.MAS.64.L.A.I | Store multiplex-and-swap octlet little-endian aligned immediate |
| S.MUX.64.B.A.I | Store multiplex octlet big-endian aligned immediate |
| S.MUX.64.L.A.I | Store multiplex octlet little-endian aligned immediate |

| size | ordering | | alignment | |
|-----------------------|----------|---|-----------|--|
| 8 | | | | |
| 16 32 64 128 | L | B | | |
| 16 32 64 128 | L | B | A | |

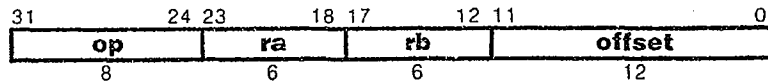
⁴⁷S.8.I need not specify byte ordering, nor need it specify alignment checking, as it stores a single byte.

Terpsichore System Architecture

Wed, Aug 2, 1995

Format

S.size.order.align.l ra,rb,offset

Description

A virtual address is computed from the sum of the contents of register ra and the sign-extended value of the offset field. The contents of register rb, treated as the size specified, is stored in memory using the specified byte order.

If alignment is specified, the computed virtual address must be aligned, that is, it must be an exact multiple of the size expressed in bytes. If the address is not aligned an "access disallowed by virtual address" exception occurs.

Definition

def StoreImmediate(op,ra,rb,offset) as

case op of

S8I,

S16LI, S16LAI, S16BI, S16BAI,

S32LI, S32LAI, S32BI, S32BAI,

S64LI, S64LAI, S64BI, S64BAI,

S128LI, S128LAI, S128BI, S128BAI:

function ← NONE

SAAS64BAI, SAAS64LAI:

function ← AAS

SCAS64BAI, SCAS64LAI:

function ← CAS

SMAS64BAI, SMAS64LAI:

function ← MAS

SMUX64BAI, SMUX64LAI:

function ← MUX

endcase

case op of

S8I:

size ← 8

S16LI, S16LAI, S16BI, S16BAI:

size ← 16

S32LI, S32LAI, S32BI, S32BAI:

size ← 32

S64LI, S64LAI, S64BI, S64BAI, SAAS64BAI, SAAS64LAI,

SCAS64BAI, SCAS64LAI, SMAS64BAI, SMAS64LAI, SMUX64BAI, SMUX64LAI:

size ← 64

S128LI, S128LAI, S128BI, S128BAI:

size ← 128

endcase

case op of

S8I,

S16LI, S16LAI, S16BI, S16BAI,

S32LI, S32LAI, S32BI, S32BAI,

S64LI, S64LAI, S64BI, S64BAI,

SAAS64BAI, SAAS64LAI:

Terpsichore System Architecture

Wed, Aug 2, 1995

```

        rsize ← 64
        SCAS64BAI, SCAS64LAI, SMAS64BAI, SMAS64LAI, SMUX64BAI, SMUX64LAI:
        rsize ← 128
        S128LI, S128LAI, S128BI, S128BAI:
        rsize ← 128
    endcase
    case op of
    S8I:
        align ← undefined
        S16LI, S32LI, S64LI, S128LI,
        S16BI, S32BI, S64BI, S128BI:
        align ← false
        S16LAI, S32LAI, S64LAI, S128LAI,
        S16BAI, S32BAI, S64BAI, S128BAI,
        SAAS64BAI, SAAS64LAI, SCAS64BAI, SCAS64LAI,
        SMAS64BAI, SMAS64LAI, SMUX64BAI, SMUX64LAI:
        align ← true
    endcase
    case op of
    S8I:
        order ← undefined
        S16LI, S32LI, S64LI, S128LI,
        S16LAI, S32LAI, S64LAI, S128LAI,
        SAAS64LAI, SCAS64LAI, SMAS64LAI, SMUX64LAI:
        order ← L
        S16BI, S32BI, S64BI, S128BI,
        S16BAI, S32BAI, S64BAI, S128BAI,
        SAAS64BAI, SCAS64BAI, SMAS64BAI, SMUX64BAI:
        order ← B
    endcase
    a ← RegRead(ra, 64)
    VirtAddr ← a + (offset150 || offset)
    if align then
        if (VirtAddr and ((size/8)-1)) ≠ 0 then
            raise AccessDisallowedByVirtualAddress
        endif
    endif
    m ← RegRead(rb, rsize)
    case function of
    NONE:
        StoreMemory(VirtAddr, size, order, msize-1..0)
    AAS:
        b ← LoadMemory(VirtAddr, size, order)
        StoreMemory(VirtAddr, size, order, m63..0+b)
        RegWrite(rb, 64, b)
    CAS:
        b ← LoadMemory(VirtAddr, size, order)
        if (b = m63..0) then
            StoreMemory(VirtAddr, size, order, m127..64)
        endif
        RegWrite(rb, 64, b)
    MAS:
        b ← LoadMemory(VirtAddr, size, order)
        n ← (m127..64 & m63..0) | (b & ~m63..0)
        StoreMemory(VirtAddr, size, order, n)
        RegWrite(rb, 64, b)
    MUX:

```

Terpsichore System Architecture

Wed, Aug 2, 1995

```

b ← LoadMemory(VirtAddr,size,order)
n ← (m127..64 & m63..0) | (b & ~m63..0)
StoreMemory(VirtAddr,size,order,n)

```

```

endcase
enddef

```

Exceptions

Reserved instruction
 Access disallowed by virtual address
 Access disallowed by tag
 Access disallowed by global TLB
 Access disallowed by local TLB
 Access detail required by tag
 Access detail required by local TLB
 Access detail required by global TLB
 Cache coherence intervention required by tag
 Cache coherence intervention required by local TLB
 Cache coherence intervention required by global TLB
 Local TLB miss
 Global TLB miss

| Zelus System Architecture | Tue, Aug 17, 1999 | Common Elements |
|---|-------------------|-----------------|
| Octet | | |
| An octet is the concatenation of 64 bits, and is the concatenation of eight bytes: | | |
| 63 | octet63..32 | 32 |
| 31 | octet31..0 | 0 |
| Hexlet | | |
| A hexlet is the concatenation of 128 bits, and is the concatenation of sixteen bytes: | | |
| 127 | hexlet127..96 | 96 |
| 95 | hexlet95..64 | 64 |
| 63 | hexlet63..32 | 32 |
| 31 | hexlet31..0 | 0 |

| Zelus System Architecture | Tue, Aug 17, 1999 | Common Elements |
|---|-------------------|-----------------|
| Triclet | | |
| A triclet is the concatenation of 256 bits, and is the concatenation of thirty-two bytes: | | |
| 255 | triclet255..224 | 224 |
| 223 | triclet223..192 | 192 |
| 191 | triclet191..160 | 160 |
| 159 | triclet159..128 | 128 |
| 127 | triclet127..96 | 96 |
| 95 | triclet95..64 | 64 |
| 63 | triclet63..32 | 32 |
| 31 | triclet31..0 | 0 |

| Zelus System Architecture | Tue, Aug 17, 1999 | Common Elements |
|--|-------------------|-----------------|
| Floating-Point Data | | |
| Zelus's floating-point formats are designed to satisfy ANSI/IEEE standard 754-1985 binary floating-point Arithmetic. Standard 754 leaves certain aspects to the discretion of implementers: additional precision, formatting of quiet and signaling NaN values, details of production and propagation of quiet NaN values. These aspects are detailed below. | | |
| Zelus adds additional half-precision and quad-precision formats to standard 754's single-precision and double-precision formats. Zelus's double-precision justifies standard 754's | | |

| Zelus System Architecture | Tue, Aug 17, 1999 | Common Elements |
|--|-------------------|-----------------|
| Floating-Point Data | | |
| Zelus's floating-point formats are designed to satisfy ANSI/IEEE standard 754-1985 binary floating-point Arithmetic. Standard 754 leaves certain aspects to the discretion of implementers: additional precision, formatting of quiet and signaling NaN values, details of production and propagation of quiet NaN values. These aspects are detailed below. | | |
| Zelus adds additional half-precision and quad-precision formats to standard 754's single-precision and double-precision formats. Zelus's double-precision justifies standard 754's | | |

- 13 -

MicroUnity

- 14 -

MicroUnity

| Zero System Architecture | Tue, Aug 17, 1999 | Common Elements |
|---|-------------------|-----------------|
| <p>precision requirements for a single-extended format, and Zero's quad-precision satisfies standard 754's precision requirements for a double-extended format.</p> <p>Each precision format employs fields labeled a (sign), e (exponent), and f (fraction) to encode values that are (1) NaN, quiet and signaling, (2) infinities, (3) $(-1)^s 2^{e-1}$, (4) normalized numbers, (5) $(-1)^s 2^{e-1} \text{bin}(1,0)$, (6) denormalized numbers, (7) $(-1)^s 2^{e-1} \text{bin}(0,0)$, and (8) zero, $(-1)^s 2^0$.</p> <p>Quiet NaN values are denoted by any sign bit value, an exponent field of all one bits, and a non-zero fraction with the most significant bit set. Quiet NaN values generated by default exception handling of standard operations have a zero sign bit, an exponent field of all one bits, a fraction field with the most significant bit set, and all other bits cleared.</p> <p>Signaling NaN values are denoted by any sign bit value, an exponent field of all one bits, and a non-zero fraction with the most significant bit cleared.</p> <p>Infinite values are denoted by any sign bit value, an exponent field of all one bits, and a zero fraction field.</p> <p>Normalized number values are denoted by any sign bit value, an exponent field that is not all one bits or all zero bits, and any fraction field value. The numeric value encoded is $(-1)^s 2^{e-1} \text{bin}(1,0)$. The bits in the fraction field are equal to the value resulting from adding all but the most significant bit of the exponent field, half 15, single 127, double 1023, and quad 16383.</p> <p>Denormalized number values are denoted by any sign bit value, an exponent field that is all zero bits, and a non-zero fraction field value. The numeric value encoded is $(-1)^s 2^{e-1} \text{bin}(0,0)$.</p> <p>Zero values are denoted by any sign bit value, and exponent field that is all zero bits, and a fraction field that is all zero bits. The numeric value encoded is $(-1)^s 2^0$. The distinction between +0 and -0 is significant in some operations.</p> <p>Half-precision Floating-point</p> <p>Zero half precision uses a format similar to standard 754's requirements, reduced to a 16-bit overall format. The format contains sufficient precision and exponent range to hold a 12-bit signed integer.</p> | | |
| | 15 | Memory |
| Zero System Architecture | Tue, Aug 17, 1999 | Common Elements |
| <p>Single-precision Floating-point</p> <p>Zero single precision satisfies standard 754's requirements for "single."</p> <p>Zero single precision satisfies standard 754's requirements for "double."</p> <p>Double-precision Floating-point</p> <p>Zero double precision satisfies standard 754's requirements for "double."</p> <p>Quad-precision Floating-point</p> <p>Zero quad precision satisfies standard 754's requirements for "double extended" but has additional fraction precision to use 128 bits.</p> | | |
| | 32 | Memory |
| | 64 | Memory |
| | 128 | Memory |

| | |
|---|---|
| <p>Zeus System Architecture</p> <p>Tue, Aug 17, 1999</p> <p>Zeus Processor</p> <p>The gateway contains two data buses within its structure, a code address and a new privilege level.</p> <div><div>63</div><div>code address</div><div>21 0</div><div>62</div><div>2</div></div> <p>The virtual memory system can be used to designate a region of memory as containing gateway. Other data may be placed within the gateway region, provided that if an attempt is made to use the additional data as a gateway, that security cannot be violated. For example, 64-bit data or stack pointers which are aligned to at least 4 bytes and are in little-endian byte order have plus, so that the privilege level cannot be raised by attempting to use the additional data as a gateway.</p> <p>User State</p> <p>The user state consists of hardware data structures that are accessible to all conventional compiled code. The Zeus user state is designed to be as regular as possible, and consists only of the general registers, the program counter, and virtual memory. There are no specialized registers for condition codes, operating modes, rounding modes, integer multiply/divide, or floating-point values.</p> <p>General Registers</p> <p>Zeus user state includes 64 general registers. All are identical; there is no dedicated zero-valued register, and there are no dedicated floating-point registers.</p> <div><div>127</div><div>0</div><div>REG[0]</div><div>REG[1]</div><div>REG[2]</div><div>REG[62]</div><div>REG[63]</div><div>128</div><div>2</div></div> <p>Some Zeus instructions have 64-bit register operands. These operands are sign-extended to 128 bits which written to the register file, and the low-order 64 bits are chosen when read from the register file.</p> <p>Definition</p> <pre>def val ← Reg[regnum, x12] case size of 64: val ← Reg[regnum, x12, 0] 128:</pre> | <p>Zeus System Architecture</p> <p>Tue, Aug 17, 1999</p> <p>Zeus Processor</p> <p>Program Counter</p> <p>The program counter contains the address of the currently executing instruction. This register is implicitly manipulated by branch instructions, and read by branch instructions that save a return address in a general register.</p> <div><div>63</div><div>ProgramCounter</div><div>2 10</div><div>62</div><div>2</div></div> <p>Privilege Level</p> <p>The privilege level register contains the privilege level of the currently executing instruction. This register is implicitly manipulated by branch instructions, and read by branch instructions that save a return address in a general register.</p> <div><div>63</div><div>PrivilegeLevel</div><div>2 10</div><div>62</div><div>2</div></div> <p>Program Counter and Privilege Level</p> <p>The program counter and privilege level may be packed into a single octet. This combined data structure is saved by the Branch Gateway instruction and restored by the Branch Down instruction.</p> <div><div>63</div><div>ProgramCounter</div><div>2 10</div><div>62</div><div>2</div></div> <p>System State</p> <p>The system state consists of the facilities not normally used by conventional compiled code. These facilities provide mechanisms to execute such code in a fully virtual environment. All system state is memory mapped, so that it can be manipulated by compiled code.</p> |
|---|---|

4/12

The `lt` relation can be used to determine the unordered condition of a single operand by comparing the operand with itself.

| MANAGEMENT code | Class | Branch | Direct | F | VALUES | COMPARE | AS | Exception if |
|--------------------|-------|--------|---------|------|--------|---------|--------|--------------|
| | | Unord | Greater | Less | Equal | Unord | Invald | |
| E | F | F | F | T | T | no | no | |
| EG | F | T | T | F | F | no | no | |
| L | F | F | F | F | F | no | no | |
| GE | F | T | F | T | T | no | no | |

Compasses

The following housing-pair comparisons are provided as instructions:

| Attribute | Code | C-value | Result if values compare as: | Exception if: |
|-----------|------|---------------|------------------------------|------------------------|
| | | Unord- ord | Greater Less | Equal Unord- ord |
| E | == | F | F | T |
| LG | > | F | T | F |
| LG | < | F | T | F |
| E | >= | F | T | F |
| EX | == | F | F | T |
| LGX | > | F | T | F |
| EX | < | F | T | F |
| EX | <= | F | T | F |
| CEX | <= | F | F | T |

The basic operations supported in hardware are floating-point add, subtract, multiply, divide, square root and conversions among floating-point formats and between floating-point and binary integer formats.

Software libraries provide ether operations required by the ANSI/IEEE floating-point standard.

2000 Progress

In addition to the basic operations, Zebra performs a variety of operations in which one or more products are furnished to each other and/or to an additional operand. The instructions include a fused multiply-add (FMA), add (FADD), subtract (FSUB), compare (FCOMP), matrix multiply (FMMUL), matrix-add (FMSCAL), and scale-add (FSCALADD).

Bounding and exceptions

rounding mode. Similarly, the instructions explicitly specify how standard exceptions (invalid operation, division by zero, overflow, underflow and inexact) are to be handled.¹

whereas no rounding is explicitly named by the instructions (e.g., *round* to nearest rounding mode). The floating-point exception signals cause the standard-specified behavior to occur, rather than a trap. When rounding is explicitly named by the instruction (N integer; Z2 real), as here, C cannot specify the specified rounding is performed, and floating-point exceptions other than those listed cause a floating-point exception trap. When X exact, or underflow is specified, all floating-point exception signals cause a floating-point exception trap, including inexact.

This technique uses the Zetas processor in executing floating-point operations with greater efficiency. When the floating-point instructions are executed, the Zetas processor can be interrupted. When the default rounding and exception handling control is specified in floating-point instructions, Zetas may safely return instructions following them, as they are guaranteed not to cause data-dependent exceptions. Similarly, floating-point instructions with the Z/N, P, or C control are guaranteed not to cause data-dependent exceptions once the operators have been executed to rule out invalid operations, division by zero, overflow or underflow of the floating-point instructions. Only floating-point instructions with X control, or which exceptions cannot be ruled out with N, Z, P, or C control need to avoid raising following instructions until the final result is generated.

ANSI/ISO standard 754:1993 specifies information to be given to trap handlers for the floating-point exception. The Zaus architecture produces a precise exception. (The program counter points to the instruction that caused the exception and all register state is preserved) from which all the required information can be produced in software, as all source operands values and the specified operation are available.

| | |
|---|--|
| <p>Zeus System Architecture Thu, Aug 17, 1999 Zeus Processor</p> <p>Each unit of a branch-conditional instruction, so that instructions at the target of the branch can be fetched in advance of the branch-conditional instruction reaching the execution pipeline. Placing the branch logic as early as possible, and as a point where the same instruction will not reduce the execution rate compared to performance. In other words, an optimizing compiler should insert the branch-logic instruction as early as possible in the basic block where the point will contain a more one other "front-end" instruction.</p> <p>Additional Load and Execute Resources</p> <p>Studies of the dynamic distribution of Zeus instructions on various benchmark suites indicate that the static frequency-based instruction classes are load instructions and execute instructions. In a high-performance Zeus implementation, it is advantageous to consider execution pipelines in which the ability to target the machine resources toward having load and execute instructions is increased.</p> <p>One of the means to increase the ability to issue execute-class instructions is to provide the means to issue two execute instructions in a single-issue string. The execution unit actually requires several different resources, so by partitioning these resources, the issue capability can be increased without increasing the number of functional units, other than the increased register file read and write ports. The partitioning favored for the initial implementation places all instructions that involve shifting and shuffling in one execution unit, and all instructions that involve multiplication, including fixed-point and floating-point multiply and add in another unit. Resources used for implementing add, subtract, and bitwise logical operations may be duplicated, being made in size compared to the shift and multiply units, or shared between the two units, as the operations have low-enough latency that two operations might be performed within a single-issue cycle. These instructions must generally be independent, except perhaps that two simple add, subtract, or bitwise logical instructions may be performed sequentially, if the resources for executing simple instructions are shared between the execution units.</p> <p>One of the means to increase the ability to issue load-class instructions is to provide the means to issue two load instructions in a single-issue string. This would generally increase the resources required of the data fetch unit and the data cache, but a compensating solution is to send the resources for the store instruction to execute the second load instruction. Thus, a single-issue string can then contain either two load instructions, or one load instruction and one store instruction, which uses the same register read ports and address computation resources as the basic Store instruction string. This capability also may be employed to provide support for unaligned load and store instructions, where a single-issue string may contain, as an alternative a single unaligned load or store instruction which uses the resources of the two load-class units in context to accomplish the unaligned memory operation.</p> <p>Result Forwarding</p> <p>When temporally adjacent instructions are executed by separate resources, the results of the first instruction must generally be forwarded directly to the resources used to execute the second instruction, where the result replaces a value which may have been fetched from a register file. Such forwarding paths use significant resources. A Zeus implementation must</p> | <p>Zeus System Architecture Thu, Aug 17, 1999 Zeus Processor</p> <p>generally provide forwarding resources to deal dependencies from earlier instructions within a string are immediately forwarded to later instructions, except between a first and second execution instruction as described above. In addition, when forwarding results from the execution units back to the data fetch unit, additional delay may be incurred.</p> |
|---|--|

7/12

8/12

Micro Unity

| Zan System Architecture | Tue, Aug 17, 1999 | Instruction Set | Zan System Architecture | Tue, Aug 17, 1999 | Instruction Set |
|---|-------------------|-----------------|-------------------------|-------------------|-----------------|
| <p>Description</p> <p>An operand size, expressed in bytes, is specified by the instruction. A virtual address is computed from the sum of the contents of register <i>rc</i> and the sign-extended value of the offset field, multiplied by the operand size. The contents of register <i>rd</i>, stored as the size specified, are written to memory using the specified byte order.</p> <p>The computed virtual address must be aligned, that is, it must be an exact multiple of the size expressed in bytes. If the address is not aligned in "access disallowed by virtual address" exception occurs.</p> <p>Definition</p> <p>dest: destination register or offset as case op of</p> <p>src: register</p> <p>src ← <i>B</i> <i>S16L, S16AL, S16OL, S16AL</i> src ← <i>L</i> <i>S32L, S32AL, S32OL, S32AL</i> src ← <i>B</i> <i>S64L, S64AL, S64OL, S64AL</i> src ← <i>B</i> <i>S128L, S128AL, S128OL, S128AL</i> src ← <i>B</i> <i>S256L, S256AL, S256OL, S256AL</i> src ← <i>B</i> <i>S512L, S512AL, S512OL, S512AL</i> src ← <i>B</i> <i>S1024L, S1024AL, S1024OL, S1024AL</i> src ← <i>B</i> <i>S2048L, S2048AL, S2048OL, S2048AL</i> src ← <i>B</i> <i>S4096L, S4096AL, S4096OL, S4096AL</i> src ← <i>B</i> <i>S8192L, S8192AL, S8192OL, S8192AL</i> src ← <i>B</i> <i>S16384L, S16384AL, S16384OL, S16384AL</i> src ← <i>B</i> <i>S32768L, S32768AL, S32768OL, S32768AL</i> src ← <i>B</i> <i>S65536L, S65536AL, S65536OL, S65536AL</i> src ← <i>B</i> <i>S131072L, S131072AL, S131072OL, S131072AL</i> src ← <i>B</i> <i>S262144L, S262144AL, S262144OL, S262144AL</i> src ← <i>B</i> <i>S524288L, S524288AL, S524288OL, S524288AL</i> src ← <i>B</i> <i>S1048576L, S1048576AL, S1048576OL, S1048576AL</i> src ← <i>B</i> <i>S2097152L, S2097152AL, S2097152OL, S2097152AL</i> src ← <i>B</i> <i>S4194304L, S4194304AL, S4194304OL, S4194304AL</i> src ← <i>B</i> <i>S8388608L, S8388608AL, S8388608OL, S8388608AL</i> src ← <i>B</i> <i>S16777216L, S16777216AL, S16777216OL, S16777216AL</i> src ← <i>B</i> <i>S33554432L, S33554432AL, S33554432OL, S33554432AL</i> src ← <i>B</i> <i>S67108864L, S67108864AL, S67108864OL, S67108864AL</i> src ← <i>B</i> <i>S134217728L, S134217728AL, S134217728OL, S134217728AL</i> src ← <i>B</i> <i>S268435456L, S268435456AL, S268435456OL, S268435456AL</i> src ← <i>B</i> <i>S536870912L, S536870912AL, S536870912OL, S536870912AL</i> src ← <i>B</i> <i>S1073741824L, S1073741824AL, S1073741824OL, S1073741824AL</i> src ← <i>B</i> <i>S2147483648L, S2147483648AL, S2147483648OL, S2147483648AL</i> src ← <i>B</i> <i>S4294967296L, S4294967296AL, S4294967296OL, S4294967296AL</i> src ← <i>B</i> <i>S8589934592L, S8589934592AL, S8589934592OL, S8589934592AL</i> src ← <i>B</i> <i>S17179869184L, S17179869184AL, S17179869184OL, S17179869184AL</i> src ← <i>B</i> <i>S34359738368L, S34359738368AL, S34359738368OL, S34359738368AL</i> src ← <i>B</i> <i>S68719476736L, S68719476736AL, S68719476736OL, S68719476736AL</i> src ← <i>B</i> <i>S137438953472L, S137438953472AL, S137438953472OL, S137438953472AL</i> src ← <i>B</i> <i>S274877907536L, S274877907536AL, S274877907536OL, S274877907536AL</i> src ← <i>B</i> <i>S549755815072L, S549755815072AL, S549755815072OL, S549755815072AL</i> src ← <i>B</i> <i>S1099511630144L, S1099511630144AL, S1099511630144OL, S1099511630144AL</i> src ← <i>B</i> <i>S2199023260288L, S2199023260288AL, S2199023260288OL, S2199023260288AL</i> src ← <i>B</i> <i>S4398046520576L, S4398046520576AL, S4398046520576OL, S4398046520576AL</i> src ← <i>B</i> <i>S8796093041152L, S8796093041152AL, S8796093041152OL, S8796093041152AL</i> src ← <i>B</i> <i>S17592186082304L, S17592186082304AL, S17592186082304OL, S17592186082304AL</i> src ← <i>B</i> <i>S35184372164608L, S35184372164608AL, S35184372164608OL, S35184372164608AL</i> src ← <i>B</i> <i>S70368744329216L, S70368744329216AL, S70368744329216OL, S70368744329216AL</i> src ← <i>B</i> <i>S140737488658432L, S140737488658432AL, S140737488658432OL, S140737488658432AL</i> src ← <i>B</i> <i>S281474977316864L, S281474977316864AL, S281474977316864OL, S281474977316864AL</i> src ← <i>B</i> <i>S562949954633728L, S562949954633728AL, S562949954633728OL, S562949954633728AL</i> src ← <i>B</i> <i>S1125899909267456L, S1125899909267456AL, S1125899909267456OL, S1125899909267456AL</i> src ← <i>B</i> <i>S2251799818534912L, S2251799818534912AL, S2251799818534912OL, S2251799818534912AL</i> src ← <i>B</i> <i>S4503599637069824L, S4503599637069824AL, S4503599637069824OL, S4503599637069824AL</i> src ← <i>B</i> <i>S9007199274139648L, S9007199274139648AL, S9007199274139648OL, S9007199274139648AL</i> src ← <i>B</i> <i>S18014398548279296L, S18014398548279296AL, S18014398548279296OL, S18014398548279296AL</i> src ← <i>B</i> <i>S36028797096558592L, S36028797096558592AL, S36028797096558592OL, S36028797096558592AL</i> src ← <i>B</i> <i>S72057594193117184L, S72057594193117184AL, S72057594193117184OL, S72057594193117184AL</i> src ← <i>B</i> <i>S144115188386234368L, S144115188386234368AL, S144115188386234368OL, S144115188386234368AL</i> src ← <i>B</i> <i>S288230376772468736L, S288230376772468736AL, S288230376772468736OL, S288230376772468736AL</i> src ← <i>B</i> <i>S576460753544937472L, S576460753544937472AL, S576460753544937472OL, S576460753544937472AL</i> src ← <i>B</i> <i>S1152921507089874944L, S1152921507089874944AL, S1152921507089874944OL, S1152921507089874944AL</i> src ← <i>B</i> <i>S2305843014179749888L, S2305843014179749888AL, S2305843014179749888OL, S2305843014179749888AL</i> src ← <i>B</i> <i>S4611686028359499776L, S4611686028359499776AL, S4611686028359499776OL, S4611686028359499776AL</i> src ← <i>B</i> <i>S9223372056718999552L, S9223372056718999552AL, S9223372056718999552OL, S9223372056718999552AL</i> src ← <i>B</i> <i>S18446744113437999104L, S18446744113437999104AL, S18446744113437999104OL, S18446744113437999104AL</i> src ← <i>B</i> <i>S36893488226875998208L, S36893488226875998208AL, S36893488226875998208OL, S36893488226875998208AL</i> src ← <i>B</i> <i>S73786976453751996416L, S73786976453751996416AL, S73786976453751996416OL, S73786976453751996416AL</i> src ← <i>B</i> <i>S147573952907503932832L, S147573952907503932832AL, S147573952907503932832OL, S147573952907503932832AL</i> src ← <i>B</i> <i>S295147905815007865664L, S295147905815007865664AL, S295147905815007865664OL, S295147905815007865664AL</i> src ← <i>B</i> <i>S590295811630015731328L, S590295811630015731328AL, S590295811630015731328OL, S590295811630015731328AL</i> src ← <i>B</i> <i>S1180591623260311426656L, S1180591623260311426656AL, S1180591623260311426656OL, S1180591623260311426656AL</i> src ← <i>B</i> <i>S2361183246520622853312L, S2361183246520622853312AL, S2361183246520622853312OL, S2361183246520622853312AL</i> src ← <i>B</i> <i>S4722366493041245706624L, S4722366493041245706624AL, S4722366493041245706624OL, S4722366493041245706624AL</i> src ← <i>B</i> <i>S9444732986082491413248L, S9444732986082491413248AL, S9444732986082491413248OL, S9444732986082491413248AL</i> src ← <i>B</i> <i>S18889465972164982826496L, S18889465972164982826496AL, S18889465972164982826496OL, S18889465972164982826496AL</i> src ← <i>B</i> <i>S37778931944329965652992L, S37778931944329965652992AL, S37778931944329965652992OL, S37778931944329965652992AL</i> src ← <i>B</i> <i>S75557863888659931305984L, S75557863888659931305984AL, S75557863888659931305984OL, S75557863888659931305984AL</i> src ← <i>B</i> <i>S1511157277773198626119968L, S1511157277773198626119968AL, S1511157277773198626119968OL, S1511157277773198626119968AL</i> src ← <i>B</i> <i>S3022314555546397252239936L, S3022314555546397252239936AL, S3022314555546397252239936OL, S3022314555546397252239936AL</i> src ← <i>B</i> <i>S6044629111092794504479872L, S6044629111092794504479872AL, S6044629111092794504479872OL, S6044629111092794504479872AL</i> src ← <i>B</i> <i>S12089258222185989008959744L, S12089258222185989008959744AL, S12089258222185989008959744OL, S12089258222185989008959744AL</i> src ← <i>B</i> <i>S24178516444371978017919488L, S24178516444371978017919488AL, S24178516444371978017919488OL, S24178516444371978017919488AL</i> src ← <i>B</i> <i>S48357032888743956035838976L, S48357032888743956035838976AL, S48357032888743956035838976OL, S48357032888743956035838976AL</i> src ← <i>B</i> <i>S96714065777487912071767952L, S96714065777487912071767952AL, S96714065777487912071767952OL, S96714065777487912071767952AL</i> src ← <i>B</i> <i>S193428131554975824143535904L, S193428131554975824143535904AL, S193428131554975824143535904OL, S193428131554975824143535904AL</i> src ← <i>B</i> <i>S386856263109951648287071808L, S386856263109951648287071808AL, S386856263109951648287071808OL, S386856263109951648287071808AL</i> src ← <i>B</i> <i>S773712526219903296574143616L, S773712526219903296574143616AL, S773712526219903296574143616OL, S773712526219903296574143616AL</i> src ← <i>B</i> <i>S1547425052439806593148287328L, S1547425052439806593148287328AL, S1547425052439806593148287328OL, S1547425052439806593148287328AL</i> src ← <i>B</i> <i>S3094850104879613186296574656L, S3094850104879613186296574656AL, S3094850104879613186296574656OL, S3094850104879613186296574656AL</i> src ← <i>B</i> <i>S618970020975922637259314912L, S618970020975922637259314912AL, S618970020975922637259314912OL, S618970020975922637259314912AL</i> src ← <i>B</i> <i>S12379400419518452745186288224L, S12379400419518452745186288224AL, S12379400419518452745186288224OL, S12379400419518452745186288224AL</i> src ← <i>B</i> <i>S24758800839036905490372576448L, S24758800839036905490372576448AL, S24758800839036905490372576448OL, S24758800839036905490372576448AL</i> src ← <i>B</i> <i>S49517601678073810980745152896L, S49517601678073810980745152896AL, S49517601678073810980745152896OL, S49517601678073810980745152896AL</i> src ← <i>B</i> <i>S99035203356147621961489305792L, S99035203356147621961489305792AL, S99035203356147621961489305792OL, S99035203356147621961489305792AL</i> src ← <i>B</i> <i>S198070406712295243922978861536L, S198070406712295243922978861536AL, S198070406712295243922978861536OL, S198070406712295243922978861536AL</i> src ← <i>B</i> <i>S396140813425910487845957723072L, S396140813425910487845957723072AL, S396140813425910487845957723072OL, S396140813425910487845957723072AL</i> src ← <i>B</i> <i>S792281626851820975691911446144L, S792281626851820975691911446144AL, S792281626851820975691911446144OL, S792281626851820975691911446144AL</i> src ← <i>B</i> <i>S1584563253703641951383822892288L, S1584563253703641951383822892288AL, S1584563253703641951383822892288OL, S1584563253703641951383822892288AL</i> src ← <i>B</i> <i>S3169126507407283902767645784576L, S3169126507407283902767645784576AL, S3169126507407283902767645784576OL, S3169126507407283902767645784576AL</i> src ← <i>B</i> <i>S6338253014814567805535291569152L, S6338253014814567805535291569152AL, S6338253014814567805535291569152OL, S6338253014814567805535291569152AL</i> src ← <i>B</i> <i>S12676506029629135611070583138304L, S12676506029629135611070583138304AL, S12676506029629135611070583138304OL, S12676506029629135611070583138304AL</i> src ← <i>B</i> <i>S25353012059258271222141166276608L, S25353012059258271222141166276608AL, S25353012059258271222141166276608OL, S25353012059258271222141166276608AL</i> src ← <i>B</i> <i>S50706024118516542444282335353216L, S50706024118516542444282335353216AL, S50706024118516542444282335353216OL, S50706024118516542444282335353216AL</i> src ← <i>B</i> <i>S101412048237033084888564670706432L, S101412048237033084888564670706432AL, S101412048237033084888564670706432OL, S101412048237033084888564670706432AL</i> src ← <i>B</i> <i>S202824096474066169777129341412864L, S202824096474066169777129341412864AL, S202824096474066169777129341412864OL, S202824096474066169777129341412864AL</i> src ← <i>B</i> <i>S405648192948132339554256682825728L, S405648192948132339554256682825728AL, S405648192948132339554256682825728OL, S405648192948132339554256682825728AL</i> src ← <i>B</i> <i>S811296385896264679108451365651456L, S811296385896264679108451365651456AL, S811296385896264679108451365651456OL, S811296385896264679108451365651456AL</i> src ← <i>B</i> <i>S162259277179252935821692731130288L, S162259277179252935821692731130288AL, S162259277179252935821692731130288OL, S162259277179252935821692731130288AL</i> src ← <i>B</i> <i>S324518554358505871643385462260576L, S324518554358505871643385462260576AL, S324518554358505871643385462260576OL, S324518554358505871643385462260576AL</i> src ← <i>B</i> <i>S64903710871710114328677092452115328L, S64903710871710114328677092452115328AL, S64903710871710114328677092452115328OL, S64903710871710114328677092452115328AL</i> src ← <i>B</i> <i>S129807421734202286573340189090230656L, S129807421734202286573340189090230656AL, S129807421734202286573340189090230656OL, S129807421734202286573340189090230656AL</i> src ← <i>B</i> <i>S25961484346840457314668037818046131328L, S25961484346840457314668037818046131328AL, S25961484346840457314668037818046131328OL, S25961484346840457314668037818046131328AL</i> src ← <i>B</i> <i>S5192296869368091462933607563609226656L, S5192296869368091462933607563609226656AL, S5192296869368091462933607563609226656OL, S5192296869368091462933607563609226656AL</i> src ← <i>B</i> <i>S10384593738736183258667215127218453312L, S10384593738736183258667215127218453312AL, S10384593738736183258667215127218453312OL, S10384593738736183258667215127218453312AL</i> src ← <i>B</i> <i>S20769187477472366517334434254436906624L, S20769187477472366517334434254436906624AL, S20769187477472366517334434254436906624OL, S20769187477472366517334434254436906624AL</i> src ← <i>B</i> <i>S415383749549447330346688685088738132512L, S415383749549447330346688685088738132512AL, S415383749549447330346688685088738132512OL, S415383749549447330346688685088738132512AL</i> src ← <i>B</i> <i>S830767499098894660693377370177476265024L, S830767499098894660693377370177476265024AL, S830767499098894660693377370177476265024OL, S830767499098894660693377370177476265024AL</i> src ← <i>B</i> <i>S16615349981977893213</i></p> | | | | | |

Kennedy Film & Book

These operations take two values from registers, perform a group of floating-point arithmetic operations on portions of bits in the operands, and place the calculated results in a register.

MicroUnity

12/12